AD-A193 059 OPTICAL MATCHED FILTERS FOR AUTONOMOUS INFRARED SEEKERS 1/2
(U) GRUMMAN CORP BETHPAGE NY SYSTEMS SCIENCES
DIRECTORATE J MENDELSOHN ET AL. 14 MAY 87 RE-735
UNCLASSIFIED NV-1-13 DAAL02-85-C-0144 F/G 17/5.1 NL

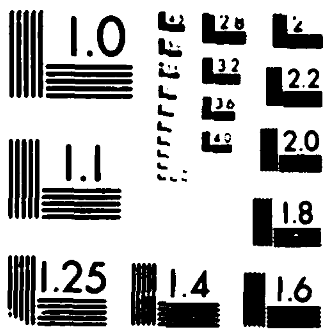

MICROCOPY RESOLUTION TEST CHART

DTIC FILE COPY

Center for Night Vision and Electro-Optics

DAAL02-85-C-0144

AD-A193 859

FINAL REPORT RE-735
OPTICAL MATCHED FILTERS FOR AUTONOMOUS INFRARED SEEKERS

MAY 1987

by
Jay Mendelsohn
and
David Englund

Systems Sciences Directorate
Grumman Corporate Research Center
Bethpage, New York 11714-3580

CONQUEST OF DARKNESS

FORT BELVOIR, VA 22060-5677

88 1 27 190

REPORT RE-735

OPTICAL MATCHED FILTERS FOR AUTONOMOUS INFRARED SEEKERS

May 1987

prepared by

Jay Mendelsohn
and
David Englund

Systems Sciences Directorate
Grumman Corporate Research Center
Bethpage, New York 11714-3580

Final Report on
Contract No. DAAL02-85-C-0144

Approved by Richard A. Scheuing
Richard A. Scheuing, V.P.
Corporate Research Center

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

Jeanie #6434

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION: Unclassified	1b RESTRICTIVE MARKINGS: None
2a SECURITY CLASSIFICATION AUTHORITY: ~~DD254~~ N/A	3 DISTRIBUTION/AVAILABILITY OF REPORT: ~~See CDRL A005~~ APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED
2b DECLASSIFICATION/DOWNGRADING SCHEDULE	
4 PERFORMING ORGANIZATION REPORT NUMBER(S): RE-735	5 MONITORING ORGANIZATION REPORT NUMBER(S)
6a NAME OF PERFORMING ORGANIZATION: Grumman Corporation; 6b OFFICE SYMBOL (If applicable)	7a NAME OF MONITORING ORGANIZATION: Center for Night Vision & Electro-Optics
6c ADDRESS (City, State, and ZIP Code): Bethpage, New York 11714-3580	7b ADDRESS (City, State, and ZIP Code): Fort Belvoir, VA 22060-5677
8a NAME OF FUNDING/SPONSORING ORGANIZATION: US Army ~~CECOM~~, CNVEO; 8b OFFICE SYMBOL (If applicable): ~~[illegible]~~ RD-NV-L	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER: Contract DAAL02-85-C-0144
8c ADDRESS (City, State and ZIP Code): Fort ~~Monmouth, NJ 07703-5000~~ Belvoir, VA 22060-5677	10 SOURCE OF FUNDING NUMBERS: PROGRAM ELEMENT NO. 62712A; PROJECT NO. 1L162712-DH95 / DH95; TASK NO. DO / 116; WORK UNIT ACCESSION NO. DA 308939

11 TITLE (Include Security Classification)

Optical Matched Filters for Autonomous INFRARED Seekers (U)

12 PERSONAL AUTHOR(S)

Jay Mendelsohn, David Englund

13a TYPE OF REPORT	13b TIME COVERED	14 DATE OF REPORT (Year, Month, Day)	15 PAGE COUNT
Final	FROM 10/85 TO 1/87	5/14/87	

16 SUPPLEMENTARY NOTATION

17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	Autonomous Seekers, Infrared, Optical Matched Filters
9	2	1	

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

The purpose of this study was to determine the capability of optical matched filtering to detect and classify infrared tactical targets. The imagery on which this determination was made was provided by the Environmental Research Institute of Michigan. The dataset consisted of 14 target types, in 370 infrared scenes of varying contrasts and intensities, from ranges of 250 m to over 12,000 m. The approach was to use a computer simulation of the optical matched filter developed by the Grumman Corporate Research Center. The Grumman approach to the problems of scale and rotational variability is to use multiple versions of the target in the reference memory. The details of this approach and the ways in which the results of the multiple reference memory are combined in the correlation plane are discussed. Matched filters were derived from the imagery as well as synthetically generated, and results for both are discussed in detail for a variety of filter-dependent and scene-

20 DISTRIBUTION/AVAILABILITY OF ABSTRACT: ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION
22a NAME OF RESPONSIBLE INDIVIDUAL	22b TELEPHONE (Include Area Code); 22c OFFICE SYMBOL

DD FORM 1473, 84 MAR — 83 APR edition may be used until exhausted. All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE
UNCLASSIFIED

dependent situations. Results are expressed in terms of detection probabilities and false alarm rates and where possible operating characteristic curves are given for various ranges and threshold values. Finally, the results point strongly to features of the approach which require improvement, and recommendations for future studies which could achieve such improvements are made.

The pen and ink changes made on the DD Form 1473 are authentic.
Per Mr. Byong H. Ahn, CNVEO/NV-RD-L

CONTENTS

CONTENTS (contd)

ILLUSTRATIONS

ILLUSTRATIONS (contd)

TABLES

TABLES (contd)

1.0 SUMMARY

The purpose of this effort was to make an initial determination of the capability of optical correlators for infrared (IR) target detection and classification, based on imagery provided to us by the Environmental Research Institute of Michigan (ERIM).

The approach throughout the study was to use a simulation of the optical filter developed by the Grumman Corporate Research Center to determine detection and classification probabilities. The basic Grumman hardware approach to the problem of size and rotational variability is to store multiple versions of the reference target and to use multiple holographic lenses to achieve, in real time, the Fourier Transform products of the incoming scene with the stored references. This method is in contrast to those that attempt to incorporate in a single filter, size and rotational invariance. In essence these methods result in a filter which is a weighted average of the reference images. After the reference images are combined there is no way of recovering the individual correlation planes that each reference image would produce when correlated against the incoming scene. The procedure which we are using allows for the combining (by addition) of the filters in the correlation plane in which case the result is equivalent to an average of the reference images with each image having weight one. It also allows, for the examination of the individual correlation planes resulting from each reference. The final decision concerning the presence of a target and the classification of that target can then be made on the basis of all of the information available for all of the correlation planes. We believe that the latter technique best preserves the information contained in the scenes and utilizes the hardware advantages of the optical approach (speed of correlation) to the maximum.

The dataset for the Characterization of Algorithms for Passive Infrared (CAPIR) study was extensive and varied. It involved 370 infrared scenes. Most scenes contained multiple target types, of which there were 14. The data was taken at two locations for which the contrast levels, background radiances and general background texture were radically different. Range variations were approximately from 250 m to 12,400 m resulting in target images which ran from well-defined shapes to small hot spots to simple points of brightness

only a few pixels in dimension. The extent of variation in the dataset and the number of targets and target views for which matched filters were required, resulted in extraordinary computational requirements for the simulation. We used a supercomputer, the Cray 1M, which is directly linked to our Image Processing Laboratory over a high speed communications line (HYPERchannel), to perform the computations. The dataset also put severe demands on the optical correlation technique in the following sense: correlators are basically template matching devices. At each position in a scene of interest, a reference shape is compared to what is in the scene. Those positions in the scene which are closely matched by the stored references are possible targets. Normally, we would expect that such a technique would work well when the targets have well defined shapes and boundaries. Infrared targets, however, as seen in the ERIM dataset, and depending on the imaging conditions and sensors, can range from having well defined silhouettes to basically having no silhouettes and being defined only in terms of blobs or hot spots. Preliminary studies (Ref 1) have indicated to us that it might be possible to use correlators on such image sets. In the case where the shape of the target was well defined, the correlator would work as expected. In the case in which the targets were hot but shapeless, enough energy could be collected by the correlators at the hot spot locations, to point to the target. In the latter case, of course, no classification could be possible. Our involvement in the CAPIR program was directed at the further investigation of these ideas in the context of a controlled and ground truthed dataset.

The principal conclusions of this study are the following:

1. Using images from the actual dataset to construct the matched filter reference memory, we were able to obtain approximately 82% detections with a false alarm rate of 1.1/scene, for targets at less than 5,200 m. These results were obtained despite the fact that a systematically placed sample of real images was not available during the course of the study. The best digital methods for which comparable results are available yielded 90% detection rates at the same level of false alarms.

2. Using synthetically generated reference images, we were able to obtain results of 78% detections with a false alarm rate of 2/scene. These results could be greatly improved using more detailed synthetic imagery.

3. In the IR imagery we studied, the matched filters were much more sensitive to intensity than they were to shape which explains the high false alarm rates that existed in some individual scenes.

4. Lack of time prevented in-depth investigation and analysis of classification capabilities but initial results are poor. This is consistent with the observation that the filters are more sensitive to intensity than to shape. Significant improvement could be possible if better filter normalization techniques are developed.

5. The principal areas that should be investigated for improving the above results are correlation plane analysis and matched filter normalizations. New techniques for carrying out these functions could result in significant improvement.

6. Results on the characterization set were not as good as the corresponding results for the digital approaches. The difference between optical and digital was much narrower on the development set. We feel the basic reason for these differences was the lack of a systematic set of images from which the matched filter reference memory could be constructed. An in-depth comparison of the results for the digital and optical approaches is presented at the end of Section 5.1.

The body of this report is devoted to describing the manner in which the conclusions above were reached and to elaborating on the results. For instance, full operating characteristic curves for various forms of the matched filters as a function of both range and threshold levels are presented throughout the report. Finally, detailed recommendations for future studies will be given in the final section.

For many years hardware development was the principal thrust of the optical correlator community. As we get closer to the realization of the correlators in hardware, systems problems must be addressed. These problems are concerned with the question of how to use the correlators to achieve specific goals. It is not enough simply to produce a correlation plane at the correlator output; we must know how that correlation plane can be best

analyzed to maximize detection and minimize false alarms. How can specific information about the amplitudes and shapes of correlation peaks be used to discriminate between true and false targets? How many versions of the targets are necessary? The more stored references we have, the more likely we are to find a target in a specific orientation. On the other hand, the more correlations that are performed the more likely that spurious or false targets will be detected. These are all problems which needed to be addressed in order to gain a better understanding of the correlator as it might operate on real world imagery in a typical target detection and classification problem.

As a result of this study, we now have a calibration point as to how well we can expect an optical correlator to perform against infrared tactical targets and some initial answers to the problems posed above. We also are now aware of specific directions to take in order to improve the correlator performance. In this report we will present detection results and show how these results are related to various target and scene parameters of interest (i.e., range, brightness, edge strength, etc). We will also give our interpretation of these results and point to ways in which we believe they can be strengthened and finally, suggest further studies that may be required to bring the use of correlators in target detection closer to reality. It should be emphasized that in many ways this study represents the first time a fully optical approach has been tried on imagery with the dramatic variations of the ERIM datasets. From this point of view, it has helped us to understand the shortcomings and need of our approach in this application as much as it has allowed others to compare this method to the digital approaches which are currently available.

2.0 INTRODUCTION

The work described herein was performed for the Center for Night Vision and Electrooptics (CNVEO) under contract #DAAL02-85-C-0144. Its purpose was to determine the detection and classification capabilities of an optical correlator against IR tactical targets. The approach was to use a digital simulation of the optical correlator. The imagery against which the test was made was provided by ERIM. As will be discussed in detail in the methodology section two different sets of matched filters were evaluated. The first set was constructed from targets extracted from the ERIM images. We refer to those as the "real" matched filters. The second set was constructed from computer generated wire frame models of the targets. In the latter case the polygon data bases for each target were provided to us by ERIM and we used our own image synthesis program to construct the targets at the ranges and orientations desired for the matched filter construction.

Almost all computations performed in this study were done on the Grumman Cray 1M. The use of a supercomputer for the simulation computations is dictated by the fact that a single correlation plane (i.e., a reference image correlated against a scene) involves the calculation of two two-dimensional Fourier transforms of a 512x512 image as well as a search of the resulting correlation plane, which is itself a 512x512 image. The number of times these operations are performed depends on the number of individual matched filters (as opposed to summed or weighted summed) that are in the reference memory.

Section 3 describes the composition of the ERIM datasets and some of the ERIM metrics associated with scenes and targets.

In Section 4 we discuss the methodology of the simulation, our matched filter construction techniques, and the composition of the reference memory for the two sets of matched filters used in this study.

The results of this study are presented and discussed in Section 5. As an example, we were able to obtain 81% detection with a false alarm rate of 1.11 false alarms per scene for targets at ranges less than 5,200 m. These results are for 88 scenes and 233 targets (14 different target types). In this section our complete results are presented. These show how both sets of matched filters fared against the ERIM development set and how the various

parameters of the scenes and targets affected these results.

Finally, in Section 6 the principal conclusions of this study are drawn and recommendations for future studies which, on the basis of this effort, appear to us to be necessary for the next step in the eventual successful use of optical matched filters for automatic target recognition.

3.0 DATASET DESCRIPTION

In order to better understand the techniques and results of this study it is necessary to be aware of the types of images and the diversity of scenes that constituted the ERIM development set. Therefore, in this section some properties of this dataset are given.

Tables 1-7 provide a description of the development dataset in terms of several of the ERIM metrics. Table 1 gives the average value of these metrics over the entire dataset. Tables 2-7 provide the range of the metric over each target. Figure 1 shows a variety of scenes and targets with the value of the ERIM metrics for the targets in those scenes. This figure will give some indication of the range of target sizes, shapes, and illumination which were contained in this dataset. The basic metrics included in the tables are:

- **RANGE**: Distance in meters from sensor to the center of the field of view. (Table 2)
- **ESR**: Edge Strength Ratio is the average squared edge strength measured by a standard edge operator, normalized to the local background intensity variance. (Table 3)
- **TIR2**: Target Interference Ratio Squared is the ratio of the squared average target-to-background intensity differences ("contrast") to the local background intensity variance. (Table 4)
- **TBIR2**: Target Background Interference Ratio is the ratio of the squared average target-to-background intensity differences ("contrast") to the product of the target and background standard deviations. (Table 5)
- **MEAN INTENSITY ON TARGET**: Average intensity (0-255 gray level value) of the pixels that make up the target within the digitized image. (Table 6)
- **PIXELS ON TARGET**: The number of pixels that make up the target within the digitized image. (Table 7)

Table 1 Average Values of Metrics

TOTAL TARGETS	378
TARGET TYPES	14
METRIC	**AVERAGE**
ESR	365.16
TIR2	30.99
TBIR2	6.09
MEAN INTENSITY	76.00
PIXELS ON TARGET	285.07

R87-3941-030(T)

Table 2 Target Range Distribution

TARGET	NUMBER OF OCCURRENCES	MIN	MAX	AVG
M-60	19	975	9793	3668
M-113	30	989	9901	4444
M-109	27	1050	11074	4251
M-151	24	1200	9316	4220
TR2	21	675	10011	3775
T-62	14	1928	9793	4288
T-55	46	975	9365	4415
BMP	9	2353	10011	5453
APC	37	953	11074	4368
BTR	34	1100	9836	4245
BRDM	23	1019	10011	4174
ZIL	31	850	11074	4039
UAZ	26	1100	9836	3954
TAB	37	997	9365	4885

R87-3941-031(T)

Table 3 Target ESR Distribution

TARGET	NUMBER OF OCCURRENCES	MIN	MAX	AVG
M-60	19	4.12	1851.95	459.85
M-113	30	4.89	625.93	162.03
M-109	27	14.95	470.60	130.62
M-151	24	7.3	975.72	295.13
TR2	21	4.64	568.78	155.36
T-62	14	4.81	1086.65	284.85
T-55	46	4.06	1488.36	201.35
BMP	9	26.84	366.79	163.69
APC	37	3.02	1342.56	175.31
BTR	34	7.14	2366.11	382.43
BRDM	23	10.51	1352.40	274.08
ZIL	31	3.39	906.35	236.59
UAZ	26	13.31	1195.40	374.31
TAB	37	9.73	7233.00	1431.79

R87-3941-032(T)

Table 4 Target TIR2 Distribution

TARGET	NUMBER OF OCCURRENCES	MIN	MAX	AVG
M-60	19	.10	180.55	46.42
M-113	30	.07	36.05	13.36
M-109	27	.30	32.23	9.26
M-151	24	.51	50.88	21.68
TR2	21	.03	41.58	12.06
T-62	14	.09	78.8	25.85
T-55	46	.01	70.21	16.74
BMP	9	.09	19.39	9.25
APC	37	.05	41.89	9.99
BTR	34	.23	175.98	34.96
BRDM	23	.00	81.42	19.00
ZIL	31	.00	78.06	24.46
UAZ	26	.23	80.78	27.26
TAB	37	.00	554.15	127.85

R87-3941-033(T)

Table 5 Target TBIR2 Distribution

TARGET	NUMBER OF OCCURRENCES	MIN	MAX	AVG
M-60	19	.09	16.08	8.11
M-113	30	.06	22.16	6.50
M-109	27	.33	12.39	3.52
M-151	24	.42	102.17	11.25
TR2	21	.03	5.71	3.15
T-62	14	.12	12.83	5.44
T-55	46	.01	14.23	5.02
BMP	9	.09	8.16	3.55
APC	37	.07	8.18	3.21
BTR	34	.30	15.90	6.01
BRDM	23	.00	7.68	3.30
ZIL	31	.00	34.00	6.90
UAZ	26	.19	30.68	7.65
TAB	37	.00	45.31	10.06

R87-3941-034(T)

Table 6 Target Mean Intensity Distribution

TARGET	NUMBER OF OCCURRENCES	MIN	MAX	AVG
M-60	19	59	118	82
M-113	30	60	82	72
M-109	27	55	88	70
M-151	24	54	89	72
TR2	21	58	79	67
T-62	14	58	105	74
T-55	46	49	98	70
BMP	9	59	72	66
APC	37	54	100	75
BTR	34	60	109	85
BRDM	23	58	81	67
ZIL	31	55	103	78
UAZ	26	58	116	84
TAB	37	56	151	93

R87-3941-035(T)

Table 7 Pixels on Target Distribution

TARGET	NUMBER OF OCCURRENCES	MIN	MAX	AVG
M-60	19	19	3339	590
M-113	30	11	1754	289
M-109	27	14	2478	425
M-151	24	4	588	104
TR2	21	17	4816	533
T-62	14	27	493	159
T-55	46	13	1810	292
BMP	9	14	238	87
APC	37	10	1665	350
BTR	34	14	1577	238
BRDM	23	10	1032	183
ZIL	31	6	2573	334
UAZ	26	6	727	142
TAB	37	13	2098	183

R87-3941-036(T)

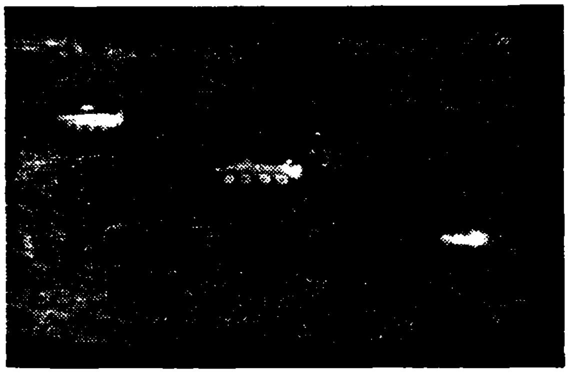

a) 1000m APC BTR UAZ

BTR METRICS

PIXELS	1577
TIR2	19.3
TBIR2	5.75
ESR	181.6
MEAN INT.	103

b) 2000m T62 M60 BTR

BTR METRICS

PIXELS	247
TIR2	47.8
TBIR2	7.22
ESR	436.8
MEAN INT.	85

c) 3000m BTR M60 T62

BTR METRICS

PIXELS	97
TIR2	90.1
TBIR2	9.09
ESR	1234.9
MEAN INT.	108

d) 5000m APC BTR UAZ

BTR METRICS

PIXELS	81
TIR2	33.04
TBIR2	9.49
ESR	532.0
MEAN INT.	87

R87-3941-001(T)

Fig. 1 Targets at Various Ranges

4.0 METHODOLOGY

The basic algorithm for the matched filter proceeds as follows: we are given a scene in which we are to determine the presence of a target. In the ERIM data bases the scenes are available in digital form (i.e., gray level (0-255) stored in raster scan form). If the scene existed on video tape or as a photograph or was coming directly from a sensor, it would first have to be digitized. Once the scene has been brought into the computer memory, its two-dimensional Fourier transform is computed and stored. This transform is then multiplied (point by point) by the matched filter stored in the reference memory. The result is then inverse Fourier transformed. The outcome of this operation is a correlation plane that is bright at those points at which the scene resembles the image from which the matched filter was made. The procedure is based on the well-known mathematical equivalence between convolution in the image domain and multiplication in the frequency domain. It should further be noted that the realization of all of these operations in the hardware device (accomplished basically using lenses and coherent light) is instantaneous.

4.0.1 Edge Filtering in the Fourier Plane

A search technique usually involving thresholding is then used to find the bright points in the correlation plane. Of course, cross correlations between items which might be close to the reference image in shape will produce bright spots whose intensities will depend on the degree of match between the reference and the object in the scene. To a certain extent, any bright area in a scene will produce some degree of cross correlation. The use of high pass matched filters provides some relief from this phenomenon. The high pass matched filters used in this study are produced by setting a disk centered at the zero frequencies equal to zero. The energy at these low frequencies is thus removed. The effect (in the image plane) is to produce an edge version of the reference. A large bright area is reduced to its edges and significant cross correlation is less likely since now the target edges and the bright area edges are less likely to be similar in shape. On the other hand, this procedure tends to reduce the total energy coming through the filter and must be applied with caution especially with respect to the actual hardware which is likely to have less dynamic range than the computer. The

procedure used to determine whether the bright spots are actually targets or merely random correlations is most important and, as we have learned from this study, should use as much a priori information about the nature of correlation peaks (expected height, shape, etc) as possible.

4.0.2 Image Source for the Construction of the Matched Filters

Two sets of matched filters were considered in this study. The first set was extracted from the ERIM development set. The second set was synthetically constructed from a polygon data base supplied to us by ERIM. The motivation for using synthetic data was that there were an insufficient number of targets in the real dataset. For example, if we wished to use an orientation resolution of 10 deg and a range resolution of 100 m for our matched filter reference memory, we would require over 5,000 versions of the targets from which matched filters would be made. In the ERIM dataset the most frequently occurring target was a T-55 which occurred only 46 times. On the other hand it seemed desirable to construct a set of matched filters based on the imagery itself since questions of how well synthetic targets actually represent real targets would, in such a case, be irrelevant. A great deal of preliminary work on both methods for constructing matched filters was carried out before arriving at the final results for these sets of matched filters. Some of this effort is of independent interest and is described in the following discussion.

4.1 REAL MATCHED FILTER METHODOLOGY

The manner in which the reference memory was constructed for the real matched filters was as follows:

We ordered the targets by the frequency of their occurrence in the development set. These were in turn ordered by range. For each target we made eight matched filters beginning with the one at the nearest range and choosing every other range. The matched filters were made by extracting the target from the scene using the ERIM truthed image as a guide. The eight matched filters were then made into a single matched filter by addition. This process was duplicated for each target. The results to be discussed are for 86 scenes containing 233 targets (all of the target types in the ERIM development set are represented in the sample).

4.1.1 Rationale for the Approach

The basic rationale behind this selection was the fact that the ERIM development set did not contain enough "close together" targets to select the matched filters in a way in which range, aspect and depression angle varied in any systematic way. We also determined in our initial studies that shape (which is governed by the three geometric variables) appeared to be less significant than intensity (hot spots) in determining the correlation plane bright spots. For this reason, we shifted the center of our matched filter plane to the target intensity center rather than its geometric center. We chose every other target so the result would not be dominated by autocorrelations and we could to some extent determine our ability to detect and classify targets not included in the matched filters. The autocorrelations on the other hand would represent how targets close to those in the matched filter set would be detected. The notion of "close" and sensitivity to shifts away from the autocorrelations could then be determined both from these results and from studies embedding synthetic targets into scenes in the development set. We fixed on range as the variable which seems to be the most important in determining the intensity distribution over the target. We summed over the matched filters for a fixed target type in order to maintain our ability to classify as well as detect. This matched filter set represents a set of targets which are embedded in the development set, and have characteristics which are representative of the overall set of targets in the development set. Finally, we did not make matched filters for target beyond certain ranges (approximately 5,500 m) since these targets appeared only as small pinpoints of light and we felt further investigation was necessary before they could be included in the matched filter set. It must be remembered that this is only the first attempt at constructing a matched filter set for a collection of targets as varied as the ones in the ERIM development set and as we shall see the matched filters performed remarkably well. Further, the experience gained from this exercise points to many areas and directions for improvement in the results.

4.1.2 Range Considerations

In Table 8 the range distributions for the targets included and excluded from the sample is given. It is clear that the included targets are biased towards the shorter ranges. As we previously discussed there appeared to be a

Table 8 Sample Statistics

	SAMPLED		NOT SAMPLED	
	#	RANGE	#	RANGE
T-55	27	986–5024	19	4924–10030
ZIL	18	751–3551	13	3508–9902
M113	16	991–3553	14	4930–9908
M109	15	1006–3560	12	3497–9997
APC	22	907–4935	15	3487–10004
BTR	19	1001–3698	15	3781–10406
UAZ	16	984–6964	7	4956–9924
M151	16	1075–5022	8	5102–10001
BRDM	17	1020–6979	6	6961–10041
TAB	20	1001–5013	17	4927–10034
TR2	16	730–6943	5	6966–9949
BMP	4	2634–5001	5	6952–9952
M60	11	882–2846	8	3762–10400
T62	7	2021–2867	7	3745–10379

R87-3941-037(T)

point where the targets appeared as only small pinpoints of light. The ramifications of this in the transform (matched filter) plane seem to be significant enough so that they should be treated separately.

4.2 WIRE FRAME MATCHED FILTER METHODOLOGY

The purpose of this portion of the study was to determine whether matched filters constructed from computer generated targets (synthetic) could be successfully used against real imagery. The advantage, of course, being that the computer generated targets can be made at arbitrary geometries and contrasts. Therefore, any matched filter density for the reference memory can be achieved. On the other hand, the problem with computer generated targets is the degree to which they are representative of the real targets and how that in turn effects the evaluation of the detection and false alarm probabilities of the matched filtering scheme.

The data for the computer generated targets was given to us by ERIM in a form which specifies vertices, edges, and surfaces. This data was used to construct the targets at a given range, aspect and depression angle using computer image generation programs developed at the Grumman Corporate Research Center. Two approaches were taken with respect to the imagery. In the first, we attempted to match the gray level shading of the computer generated target to the gray levels of the corresponding real target. In the second, only the wire frame model (with hidden lines removed) was used. From these computer generated images matched filters were made and then correlated against the scenes.

4.2.1 Ground Truth Using Correlation Analysis

Our initial approach was to perform correlation analysis on the real and synthetic targets (with gray level shading) both on a black background and in the appropriate scene. We had previously found major differences in the correlation response of the synthetic and real T-55 targets. At that time, we believed that we had generated the synthetic target at an incorrect aspect angle. This belief was confirmed when ERIM issued a correction of the aspect angle for the target and scene in question. We decided to do a correlation analysis over a large enough total change in aspect and a small enough delta change to encompass the true target aspect and determine if correlation analysis was by itself powerful enough to determine the true aspect. Thus, we

generated synthetic targets from 132 to 152 deg in aspect in steps of 2 deg and did correlation analysis on each scene. In Tables 9 and 10 the results of this analysis are shown. Table 9 shows the results of the correlations done on a black background and Table 10 for the synthetic target reinserted in the original scene. It is clear that the original ERIM estimate of the aspect of 134 deg is less satisfactory (judging from the correlation results) than the new estimate of 144 deg. However, the correlation results indicate that a still larger return is obtained at 148 deg. This in turn seems to be contradicted by the fact that visually 144 deg appears more correct. Since the ERIM matching approach depends to some extent on an operator making a visual match between the wire frame model and the target in the scene we would expect their estimate to be near the best visual match. Whether 148 deg, which is the best correlation match, is actually better than 144 deg cannot be answered here, but we believe that this analysis has shown correlation to be a potent tool for ground truthing real data and for matching synthetic data to real data.

4.2.2 Differences between Synthetic and Real Autocorrelations

It will be noted that when the synthetic target is matched against itself the maximum correlation is significantly lower than when the real target is correlated against itself. This is true at what appears to be the optimum matching aspect as well as all others. This difference appears to us to be due to the fact that we do not generate in the synthetic target the high frequency components that are present in the real targets. One reason for this is that the ERIM models consist of relatively few surfaces and since in our gray level shading scheme each surface is uniformly shaded, the high frequency content of the synthetic image tends to be decreased as opposed to the real image.

4.2.3 Correlation Sensitivity to Scale and Rotations

In Fig. 2, the results of using a target matched at one set of geometries and then rotated to another are shown. It is clear that the gray level match of 2b (synthetic) to 2a (real), for example, no longer is valid when the target in 2b is rotated to the geometry of the same target in 2c or 2e. Thus, a new shaded image is needed at every new geometric setting rather than having a single image that can be scaled and rotated to the correct geometry.

In order to use the synthetic data effectively, we chose a target and

Table 9 Correlation Results for Synthetic Targets (T-55 on Black Background)

ASPECT	SS*	RS*	SR*	RR*
132	200	74	74	354
134 (2)	208	74	74	354
136	217	80	80	354
138	219	86	86	354
140	234	101	101	354
142	225	96	96	354
144 (1)	225	104	104	354
146	224	107	107	354
148	224	109	109	354
150	221	102	102	354
152	219	97	97	354

*SS = SYNTHETIC MF - SYNTHETIC TARGET
RS = REAL MF - SYNTHETIC TARGET
SR = SYNTHETIC MF - REAL TARGET
RR = REAL MF - REAL TARGET

(1) ERIM CORRECTED ASPECT ESTIMATE
(2) ERIM INITIAL ASPECT ESTIMATE

R87-3941-038(T)

Table 10 Correlation Results for Synthetic Targets (T-55 in Scene)

ASPECT	SS*	RS*	SR*	RR*	WR*
132	192	102(X)	81	366	588
134(2)	199	102(X)	81	366	646
136	207	102(X)	83	366	70
138	210	102(X)	89	366	766
140	225	106	107	366	709
142	216	103	98	366	536
144(1)	214	111	107	366	470
146	214	114	109	366	521
148	212	117	112	366	418
150	208	109	106	366	394
152	206	104	101	366	372

*SS = SYNTHETIC MF-SYNTHETIC TARGET
RS = REAL MF-SYNTHETIC TARGET
SR = SYNTHETIC MF-REAL TARGET
RR = REAL MF-REAL TARGET
WR = WIRE FRAME MF-REAL TARGET

(1) ERIM CORRECTED ASPECT ESTIMATE
(2) ERIM INITIAL ASPECT ESTIMATE
(X) VALUE FOUND ON WRONG TARGET

R87-3941-039(T)

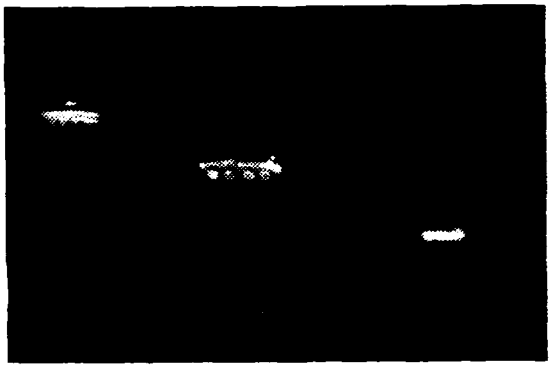

a) ORIGINAL SCENE 2.7

APC BTR UAZ

b) SYNTHETIC BTR INSERTED

c) ORIGINAL SCENE 4.1

APC BTR UAZ

d) SYNTHETIC BTR INSERTED

2 7 INTENSITIES 4.1 GEOMETRY

e) ORIGINAL SCENE 2.38

M-151 M-109 BTR

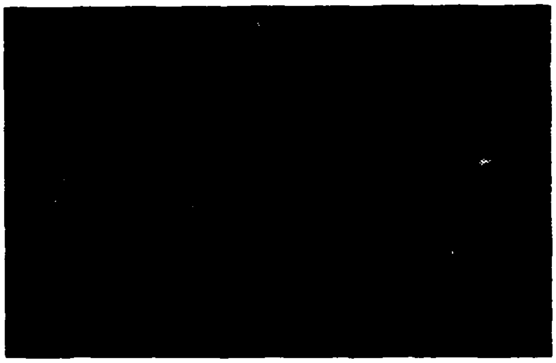

f) SYNTHETIC BTR INSERTED

2.7 INTENSITIES 2.38 GEOMETRY

R87-3941-002(T)

Fig. 2 Inserting Synthetic Targets in Real Scenes

scene in which we received one of the larger correlation peaks. This large peak was a result of the fact that the target was large and bright and had a well-defined shape. The idea was that in such a situation the correlation was more likely governed by the shape of the target than merely by the presence of hot spots on it. By varying this synthetic target in small steps in the three geometric parameters, we could get some idea of the falloff in the correlation peak and how it would depend on the shape changes that take place due to the changing geometry. We chose this approach because we observed many cases in the ERIM development set in which matched filters made from one kind of target detected other targets simply by matching hot spots, rather than truly matching shapes. Such a match, while yielding detections, carries no (and sometimes erroneous) classification information.

Figures 3, 4, and 5 show the calculated falloff for the three variables of interest: range, aspect, and depression. These were varied individually from a fixed starting position (scene 2.7-BTR target). The variations were carried out for a single variable at a time, i.e., while range was varied, aspect and depression were fixed at their starting values.

4.2.4 Correlation Sensitivity to Gray Level

In order to measure the sensitivity to changes in gray levels we correlated a matched filter made from a synthetic BTR against the same synthetic target inserted in scene 2.7. We varied the target's gray levels for -50% to +50%. The correlation values over these cases is shown in Fig. 6. The loss in the correlation as the target gray levels get smaller is close to linear. The peak followed by a decline in the correlation is probably due to the fact that, when the target gray levels saturate, the interior edges are lost. The results do indicate that the matched filters should be made so that the silhouettes obtained from the high pass versions are at maximum gray. These then would compensate to some extent for the range of gray level variations that would occur on the targets in the scenes.

4.2.5 Correlation Analysis Using Wire Frame Synthetic Models

In light of the problems associated with the shaded synthetic imagery, we decided to shift our attention to the unshaded wire frame models. It seems conceivable that using the silhouette (with hidden lines removed) might eliminate some of the problems that arise due to inadequate gray level

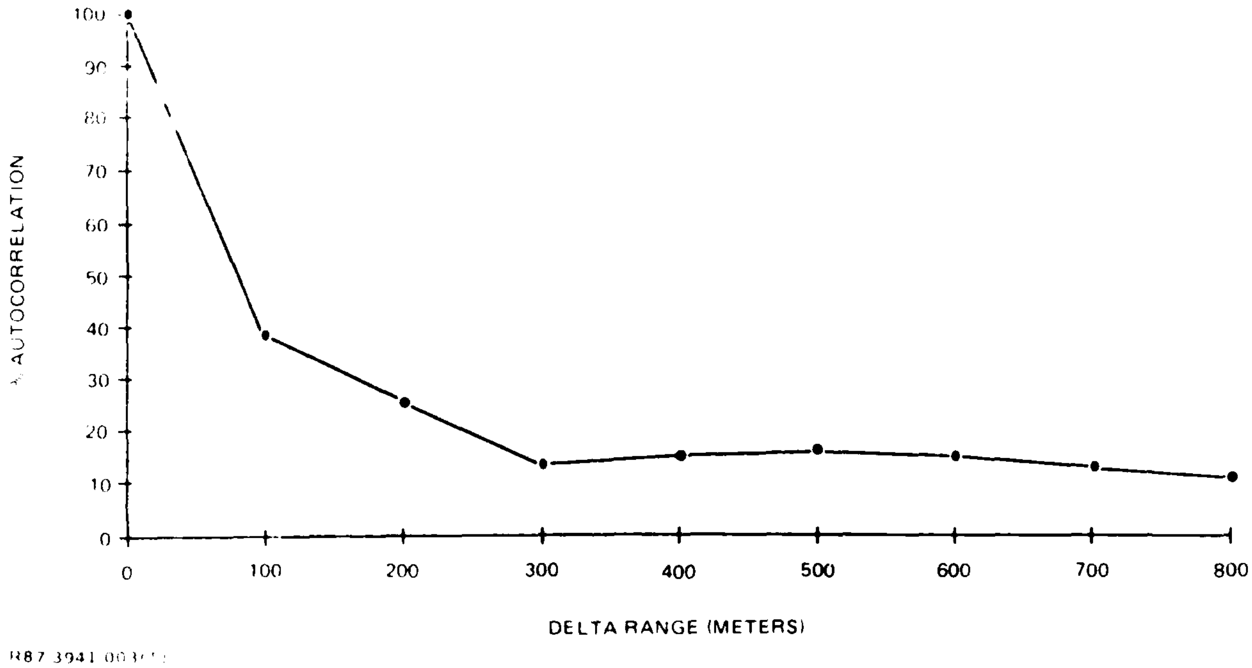

Fig. 3 Range Sensitivity

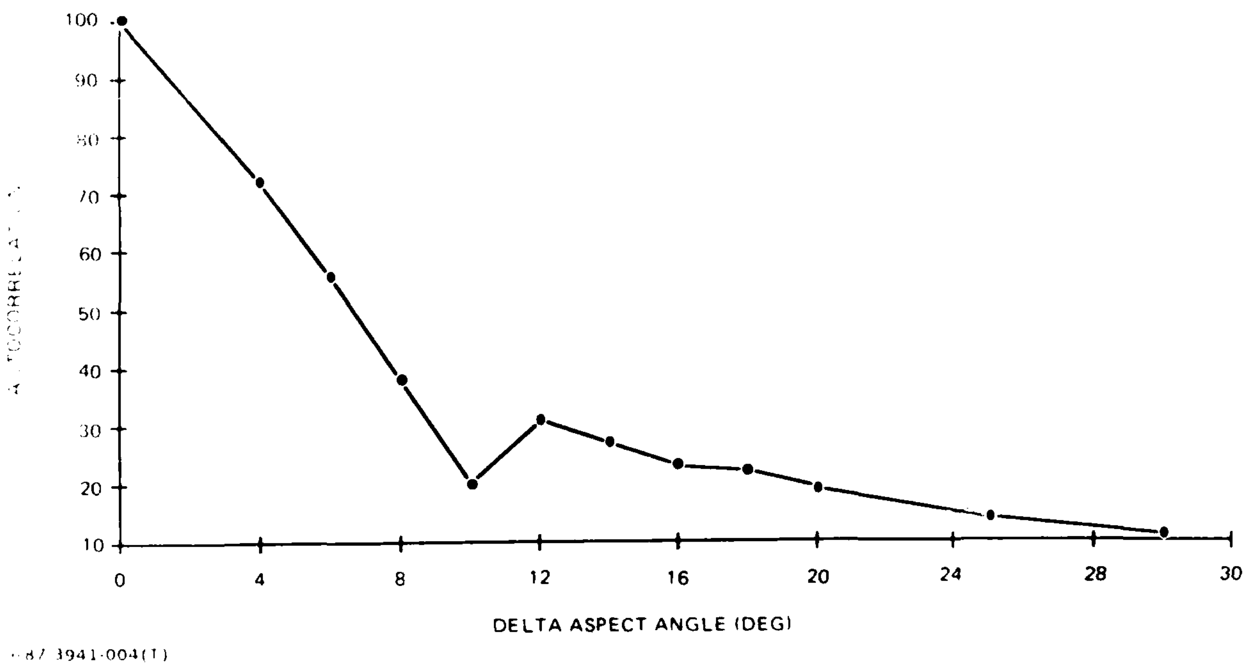

Fig. 4 Rotational Sensitivity

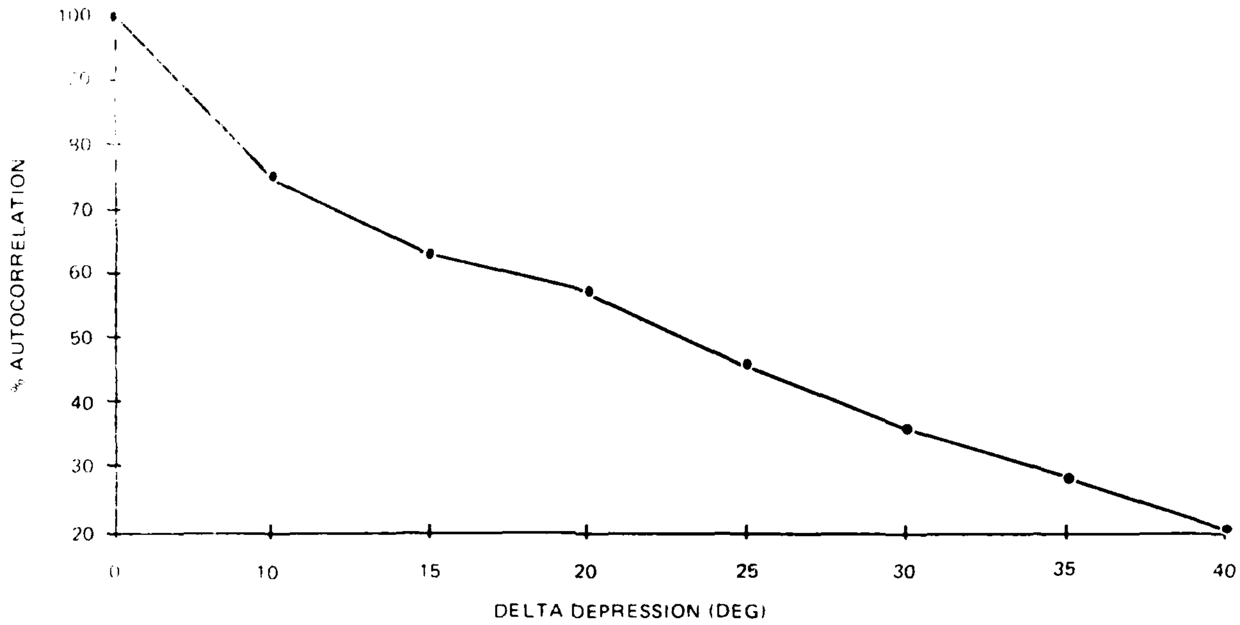

R87-3941-C05(T)

Fig. 5 Depression Angle Sensitivity

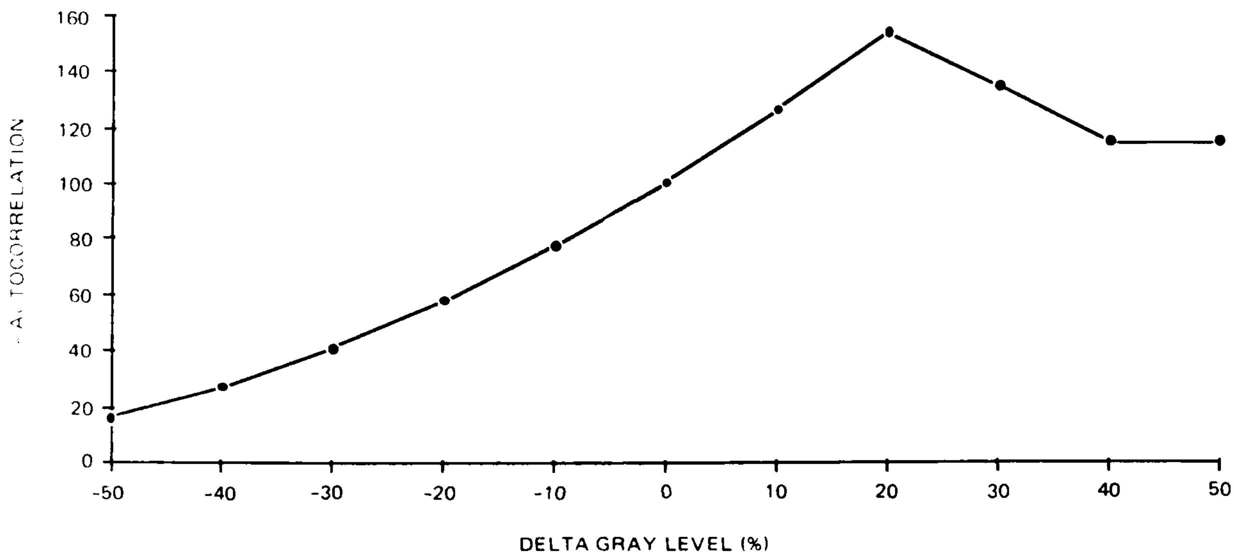

R87 3941-006(T)

Fig. 6 Gray Level Sensitivity

matches. The purpose of the high pass filter that is superimposed on the matched filter is to achieve an edge version of the image from which the matched filter is constructed. Therefore, beginning with an edge version seems justified.

In the last column of Table 10, we calculated the correlation peaks between the wire frame models (no shading) and the corresponding real image. Note from Table 10 that in absolute values the wire frame correlations return higher values than any other combination. We would expect that if the level of high pass filtering that we use in the matched filter was suitably adjusted we could match the wire frame results, since as we have just discussed the purpose of the high pass filtering is to produce a silhouette or edge version of the target much like the wire frame model generated by ERIM. We also note that the scene which yielded the maximum return was at a different aspect angle than that of the matched filters constructed from the filled polygons, and that the wire frame model results are more variable than for the filled polygons. To further examine the question of the wire frame models, we constructed a summed matched filter for which the individual matched filters were wire frame models of the T-55. We correlated this filter against 16 scenes containing T-55 targets and compared the results using the actual targets and matched filters. The wire frame models were used with both normalization with respect to the total energy in the filter and without such normalization.

The unnormalized wire frame filters were close to the reals in detections and misses but differed in the number of false alarms. Since the real filters were normalized with respect to energy, there appear to be some fundamental differences in the way the wire frames compare to the reals. However, these may be due to difference in the energy levels between the wire frame and the real matched filters. Normalization questions arose throughout the study and the resolution of these issues offers a fertile area for future research.

4.2.6 Comparisons of Synthetic Correlations to Real Correlations

In order to make further comparisons we compared the computed sensitivity of the three different synthetic models with the real model sensitivity. Results are summarized in Fig. 7. The difference between 2D synthetic and 3D synthetic is that the 2D synthetic represents a rotation of the entire plane in which the target is contained while the 3D synthetic represent a true

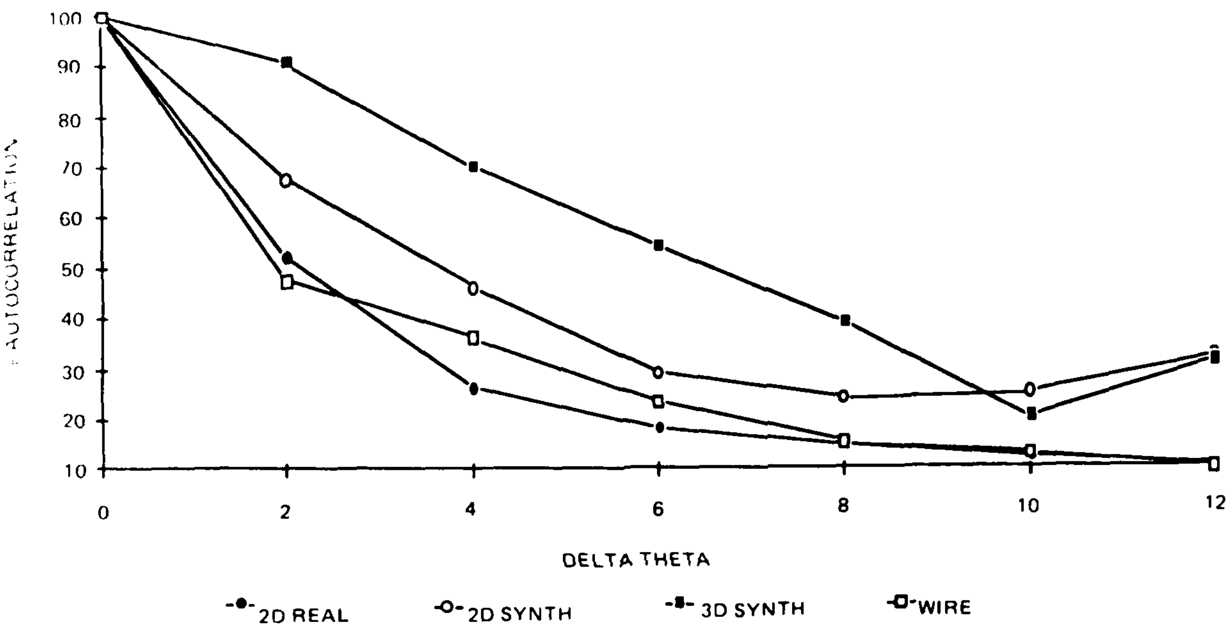

R87-3941-007(T)

Fig. 7 BTR Rotational Sensitivity

rotation of the solid model in the scene.

In the case of the 2D rotation the true target is rotated as the entire scene rotates. The values for the real sensitivity were obtained in the same fashion. In the 3D rotation the synthetic target is rotated and the scene remains constant. In rotating the synthetic target the shading of the rotated target is determined from the shading of the original target. Figure 7 clearly shows that the 3D rotation is the least comparable to the real sensitivity and that the wire frame is closest. Figure 7 appears to us to be intuitively correct in that we believe that the fundamental problem in using the synthetic images is to obtain the correct shading. Neither the wire frame nor the 2D rotations depend on such shading. The fact that the wire frame sensitivity curve is closer to the real than is the 2D synthetic may indicate that the wire frame model is a better one than the filled polygons for use in these studies.

In still another effort to determine the suitability of the wire frame models we made a matched filter from the wire frame model but varied its geometry in each variable from 1% to 5% of those values which would represent an autocorrelation with the target in the scene. The results are contained in Table 11. The target in this case was the same BTR we used for the sensitivity studies. In order to provide a frame of reference the values for the autocorrelation of the real target with itself is given. First, we notice that contrary to the T-55 results, the wire frame correlations are smaller than the real correlations. There also appears to be some question as to the value of the geometric parameters for the autocorrelation since there are two correlation peaks (1% and 3%) which are higher than those that occur at the range, aspect and depression provided in the ERIM header. While differences in magnitude need to be explained we believe that they are not as significant as the results of the second row in Table 11. This gives the result of the filter on the other target (APC) in the scene. Thus while in the real case the alternate target peak is less than 2% of the autocorrelation peak in the case of the wire frame filters the alternate targets are substantially higher percentages of the peaks on the BTR. This again shows a fundamental difference in how the synthetic targets behave. This appears to be especially true when the correlations are either between real targets and synthetic matched filters or between synthetic targets and real matched filters. It is

Table 11 Wire Frame Correlations — BTR 2.7

	REAL FILTER	% CHANGE 0	1	2	3	4	5
BTR	2629	740	1515	993	1483	878	982
APC	49	862	997	1098	993	676	628

R87-3941-040(T)

necessary to determine how far apart the matched filters should be and how many can be added together to reduce the overall number of filters that need to be correlated against a single scene. To obtain this information we investigated the auto- and cross-correlation behavior of the wire frame matched filters corresponding to the BTR target in scenes 4.1 and 3.7 (Fig. 8). In this study aspect and range was varied and depression angle held fixed at the ERIM estimate.

4.2.7 Normalization of Energy in the Matched Filter

Table 12 shows the results for the wire frame filter of the BTR in scene 4.1 at increments of 2.5 deg. Two regions appear to stand out. The first is at aspects 342.5-345 and ranges 950-1050. The other is at aspects of 337.5-340 and ranges 850-950. The ERIM estimate for this target is 342.84 in aspect and 1078 in range which agree very well with the first region of high correlation plane returns. A possible explanation for the other region is that the smaller the range the larger the target and, therefore, the larger the amount of energy in the matched filter of that target. Thus some of the resulting correlation plane energy arises merely from the energy in the filter rather than the match between the filter and the target in the scene. We tried some elementary attempts at normalization of this factor, but they did not work well enough to implement. We also should point out that, in order to get the correlations to occur on the correct target, it was necessary to apply high pass filtering to the scene. We did this at the same levels of filtering as we had applied to the filters that were made from the development set. This may not be the best value for the wire frame models and some future investigation of this factor is required.

4.2.8 Analysis of Scenes without Targets

As a basis for comparison, we correlated the same wire frame filter against one of the scenes from which we had removed the targets and these results are in Table 13. As can clearly be seen targets have very much larger peak returns than background and this remains a very positive factor for using matched filters in the detection scenario.

We further investigated the cross correlations between the wire frame filters from aspects in scene 4.1 and scene 3.7 (Table 14). These show reductions which would be sufficient to distinguish between the aspects of the

a) SCENE 4.1

b) SCENE 3.7

R87-3941-008(T)

Fig. 8 Scenes Used in Wire Frame Correlations of BTR

Table 12 4.1 Aspects vs Scene 4.1

RANGE	ASPECT 337.5	340	342.5	345	347.5	350	352.5
[illegible]	2320	2306	2231	1635	1686	1501	1538
900	2179	2393	2219	1978	1878	1929	1522
950	2866	2638	2547	2085	1453	1416	1258
1000	1996	2043	2540	2113	1765	2647	1475
1050	1720	1764	2356	2327	1749	2327	1444
1100	1500	1357	1575	1648	2248	1986	1093
1150	2158	1714	1160	1823	1876	1648	1165
1200	1520	1336	2071	1731	1542	1194	1091
1250	1979	2132	1330	1371	1210	988	856

R87-041-041(T)

Table 13 4.1 Aspects vs No Target Scene

RANGE \ ASPECT	337.5	340	342.5	345	347.5	350	352.5
850	58	67	64	79	98	81	66
900	68	57	68	88	71	55	65
950	86	59	67	59	59	46	114
1000	51	46	49	61	104	98	98
1050	61	47	58	84	92	84	64
1100	43	44	40	48	62	44	57
1150	45	42	33	47	39	37	98
1200	39	42	38	40	49	37	98
1250	33	42	38	42	60	56	40

R87 3941 042(T)

Table 14 4 1 Aspects vs Scene 3.7

RANGE	ASPECT 337.5	340	342.5	345	347.5	350	352.5
850	793	783	1073	595	607	590	580
900	619	929	606	785	654	627	684
950	1087	857	1421	897	747	749	797
1000	664	584	1079	595	873	718	594
1050	695	772	569	834	822	834	784
1100	668	916	616	829	753	811	645
1150	967	773	859	592	574	540	531
1200	1013	862	712	596	463	665	599
1250	913	665	740	531	866	494	452

R87 3941 043(T)

BTR in the two scenes.

4.2.9 Preliminary Conclusions

These preliminary results indicate that single wire frame matched filters can detect and discriminate between targets and background and between targets at grossly different aspects when the filter is applied individually and when they are densely packed (2.5 deg separation). We then pursued, on the same limited dataset, the question of how far apart the filters can get and maintain this ability. A final determination of 10 deg of rotation and 100 m of range was made on the basis of that analysis as well as on the basis of computational feasibility.

5.0 RESULTS

In the previous section we discussed our procedure for building the matched filter reference memories. In this section we will discuss how these matched filters performed against the ERIM imagery and to the best of our current understanding, the reasons for the level of performance that was achieved in each case.

5.1 REAL MATCHED FILTERS

5.1.1 <u>Methods of Correlation Plane Analysis</u>

Table 15 summarizes the most important results for this set of matched filters. Using three different techniques of correlation plane analysis, we give the detection rates (Section 5.1.2), the number of cues per scene (Section 5.1.3) and the number of false alarms per scene (Section 5.1.4).

5.1.2 <u>Simple Thresholding</u>

The first technique simply thresholds the correlation plane results obtained from the various matched filters, collects the results and reports the locations of the correlation plane peaks. This technique reflects how the correlator could be used as a "cueing" device which would point to interesting areas in the image plane and then those areas could be analyzed further.

5.1.3 <u>Correlation Peak Shape Analysis</u>

The second technique further examines the sequence of correlation plane peaks to determine if they are really peaked or merely bright. We used a comparison of the local variance and the global variance to make this determination. The results show a falloff in the detection rate but also a reduction in the false alarm rate to less than one false alarm per scene. This method puts the most stringent requirements on the correlation and, therefore, is most sensitive to how the matched filter collection actually represents the data base We believe that with improvements in the matched filter collection and with this method of correlation plane analysis we can improve the detection probability and maintain the low false alarm rate.

5.1.4 <u>Image Plane Confirmation</u>

In the third technique, we used a local image plane calculation to make

Table 15 Cueing Results

	CP INT	CP INT + SHAPE	CP INT + IMAGE
% DET	81	65	81
# OF CUES/SCENE	9.32	9.32	9.32
# OF FALSE ALARMS/SCENE	7.13	0.76	1.11

R87-3041-044(T)

the final decision concerning the absence or presence of a target at the peak locations indicated in the correlation plane. This technique maintains the original detection probability but has a higher false alarm rate than method 2.

5.1.5 Threshold Sensitivity

Figures 9-11 are graphic representations of the behavior of detections and false alarms as the correlation plane threshold is varied. Note that for method 3 there are no false alarms at gray levels about 45. This value is slightly higher for method 1. The significance of this is that there are intensity values in the correlation plane which determine a target without ambiguity.

5.1.6 Range Sensitivity

Figures 12-13 show the range sensitivity of methods 1 and 3. While there is a definite falloff in detections over range, which appears to be nearly linear, the false alarms are remarkably constant over a wide set of ranges. Qualitatively both correlation plane analysis methods are similar; however, the image plane technique is significantly better quantitatively.

5.1.7 Distribution of False Alarms

Figures 14-19 are the distribution and density functions of gray levels for detections and false alarms. The distribution functions clearly indicate again that there is a threshold value in the correlation plane such that, for gray levels above that threshold, there are no false alarms and targets are detected unambiguously. The density functions on the other hand show the areas of confusion between false alarms and detections. It is very interesting to observe the distribution of false alarms. While the absolute rate appears high the distribution of false alarms by scenes, which is given in Fig. 20-21 is quite interesting. Even for method 1 there is a significant number of scenes in which no false alarms occur. The high rate of false alarms found in method 1 is due to relatively few scenes with a very high number of false alarms. It is possible that by examining the characteristics of these scenes some common factor causing the false alarms could be determined. The image plane technique, Fig. 21, has a very large number of scenes with no false alarms and the distribution over the complete set is quite reasonable.

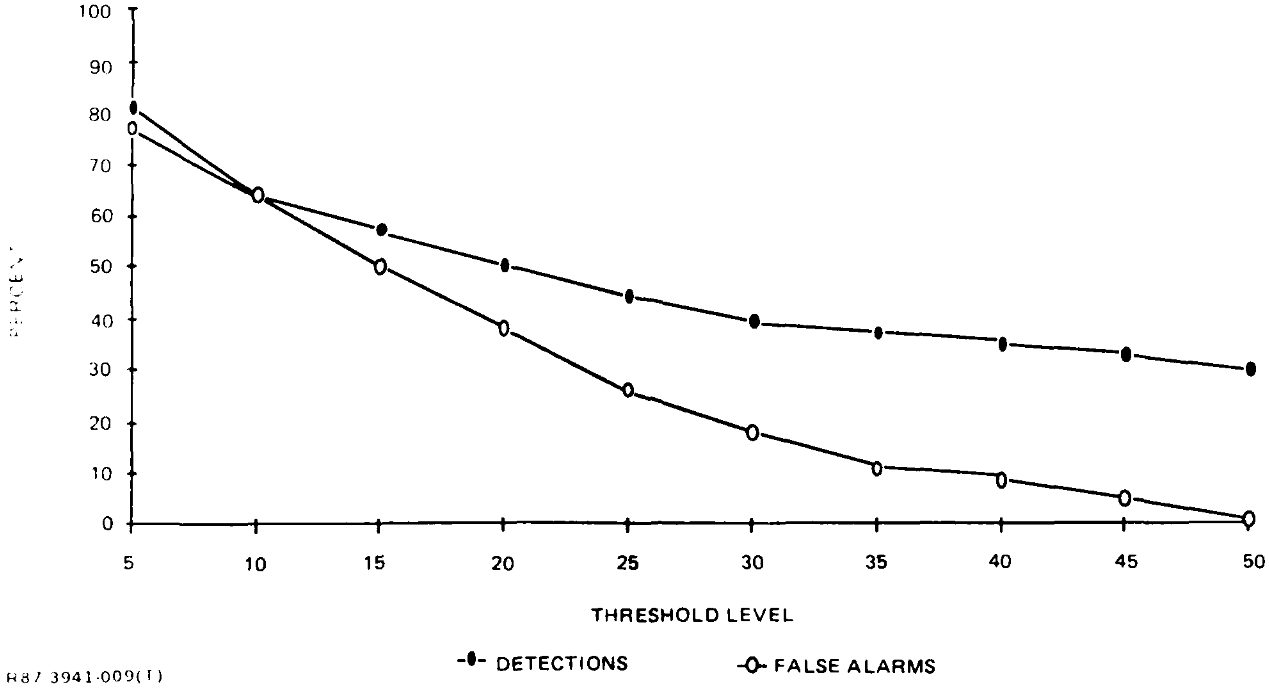

Fig. 9 Detection vs Threshold (CP Intensity)

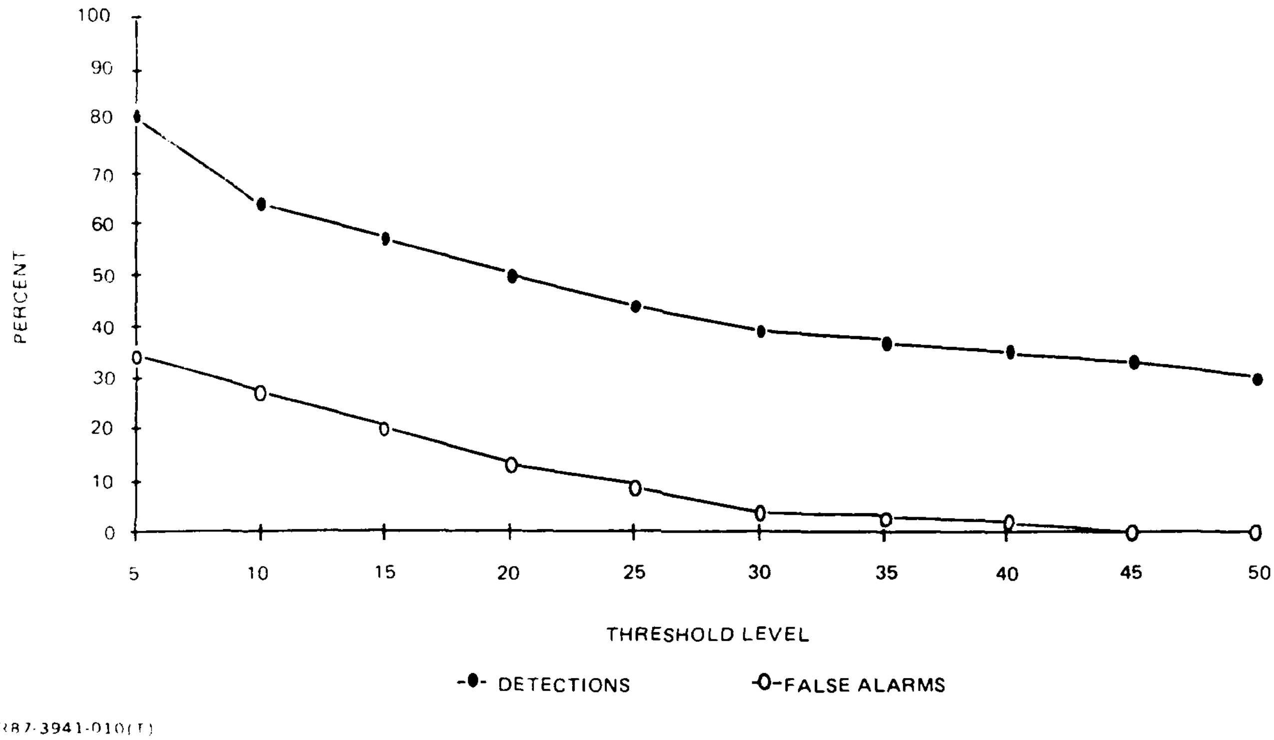

Fig. 10 Detection vs Threshold (CP Image Plane)

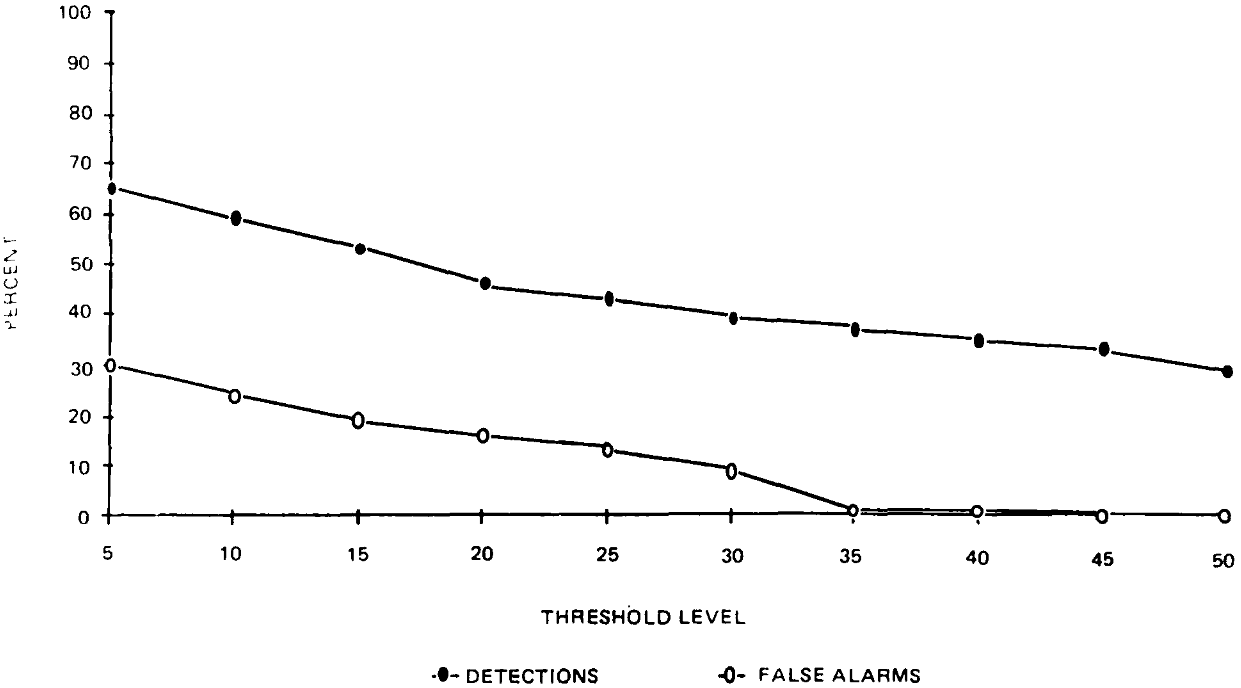

R87-3941-011(T)

Fig. 11 Detections vs Threshold (CP Shape)

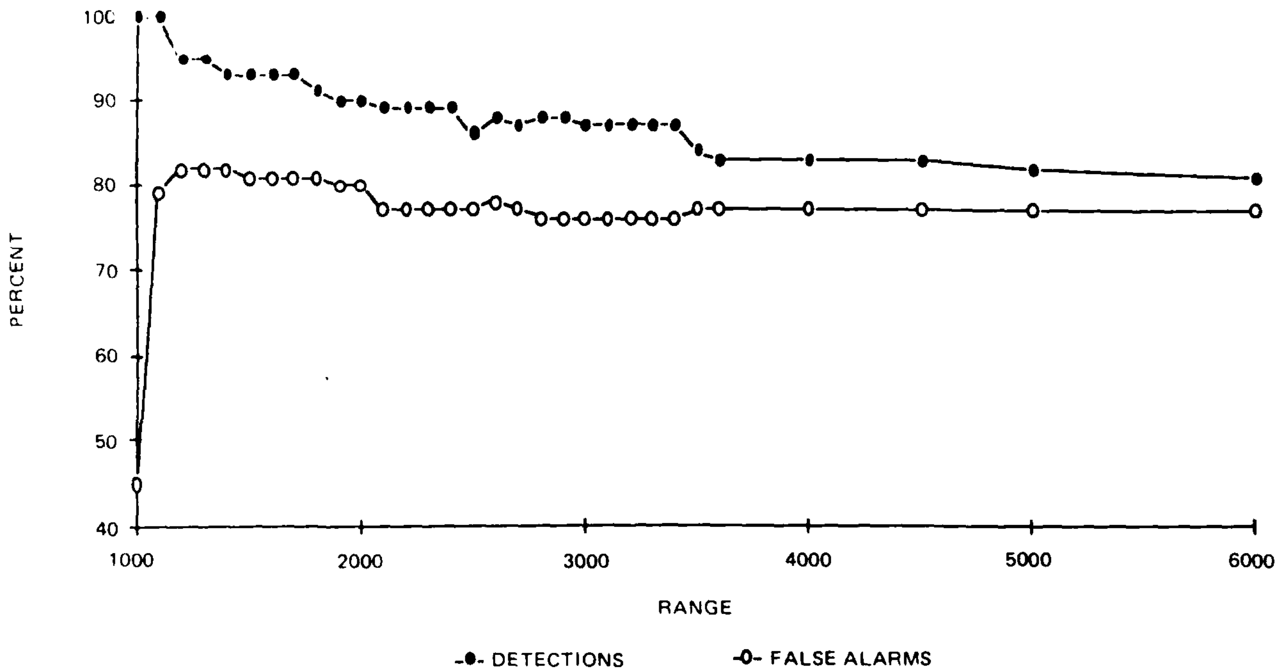

H87-3941-012(T)

Fig. 12 Range vs Detections & False Alarms

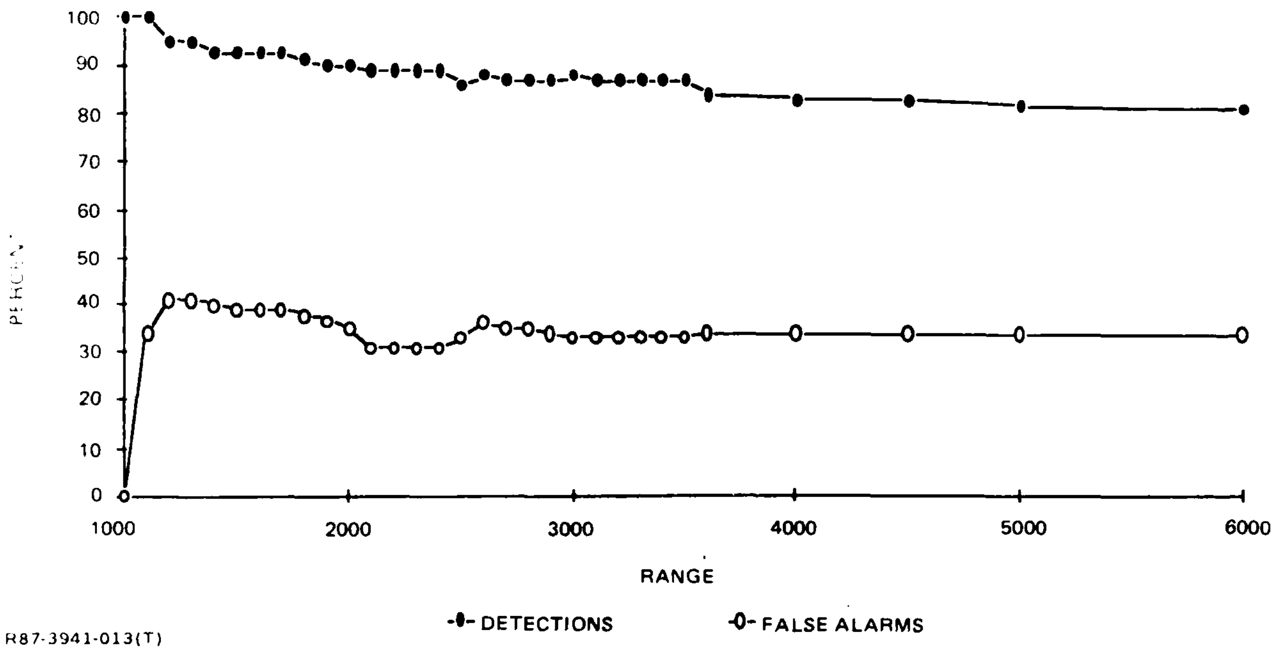

R87-3941-013(T)

Fig. 13 Range vs Detections & False Alarms (CP Image Plane)

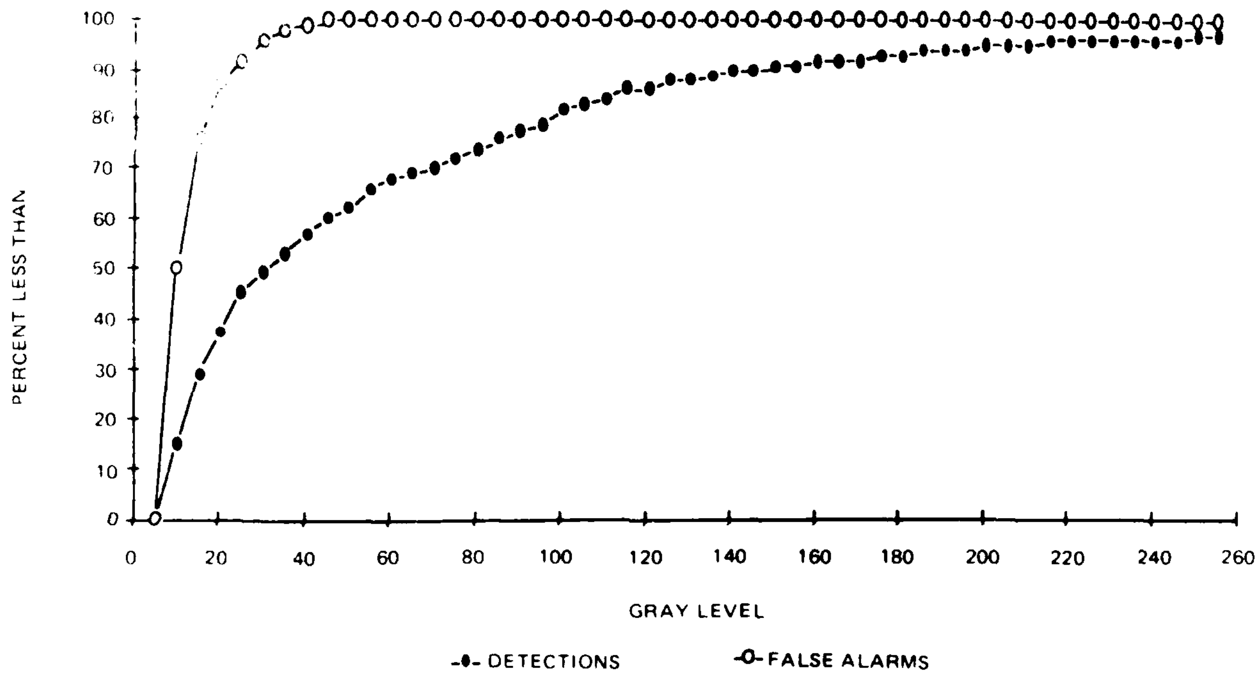

R87-3941-014(T)

Fig. 14 Cumulative Distribution of Gray Levels (CP Intensity)

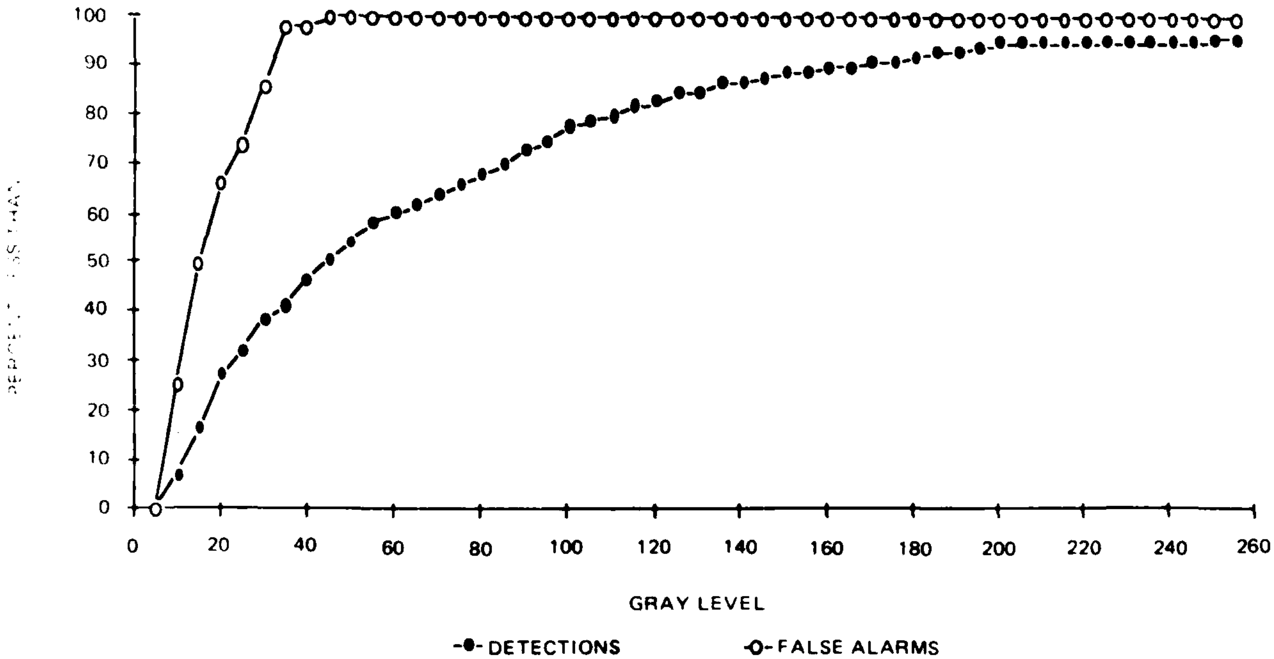

R87 3941-015(1)

Fig. 15 Cumulative Distribution of Gray Levels (CP Shape)

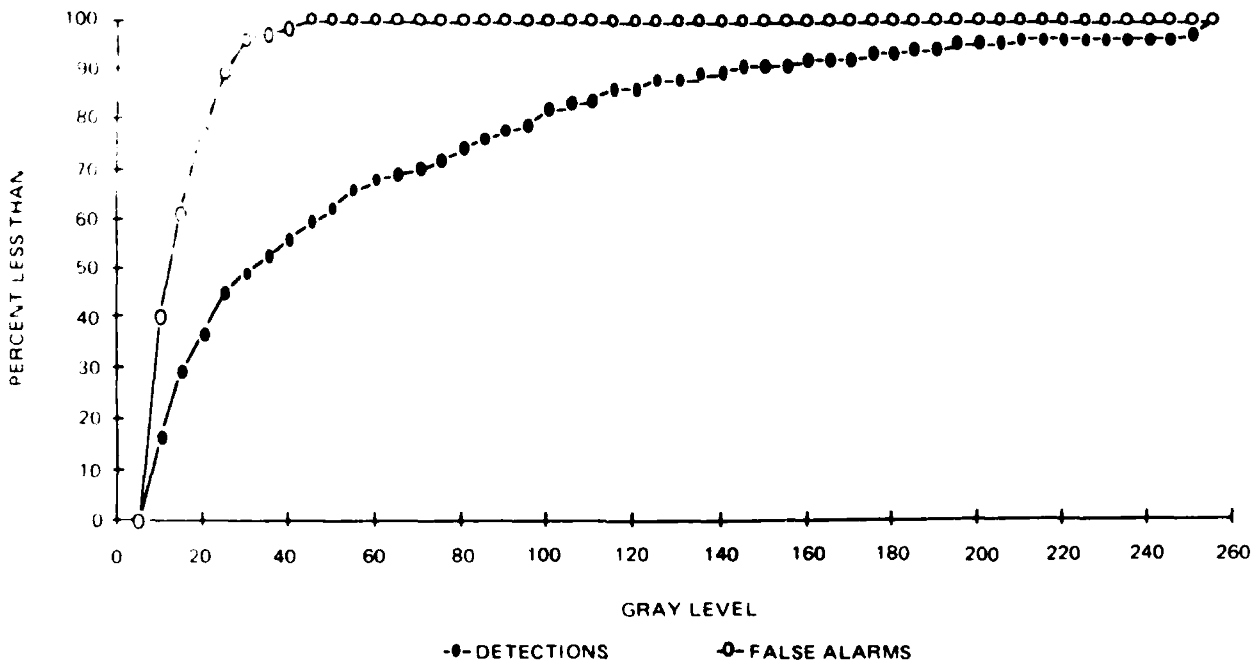

R87 3941 016(1)

Fig. 16 Cumulative Distribution of Gray Levels (CP Image Plane)

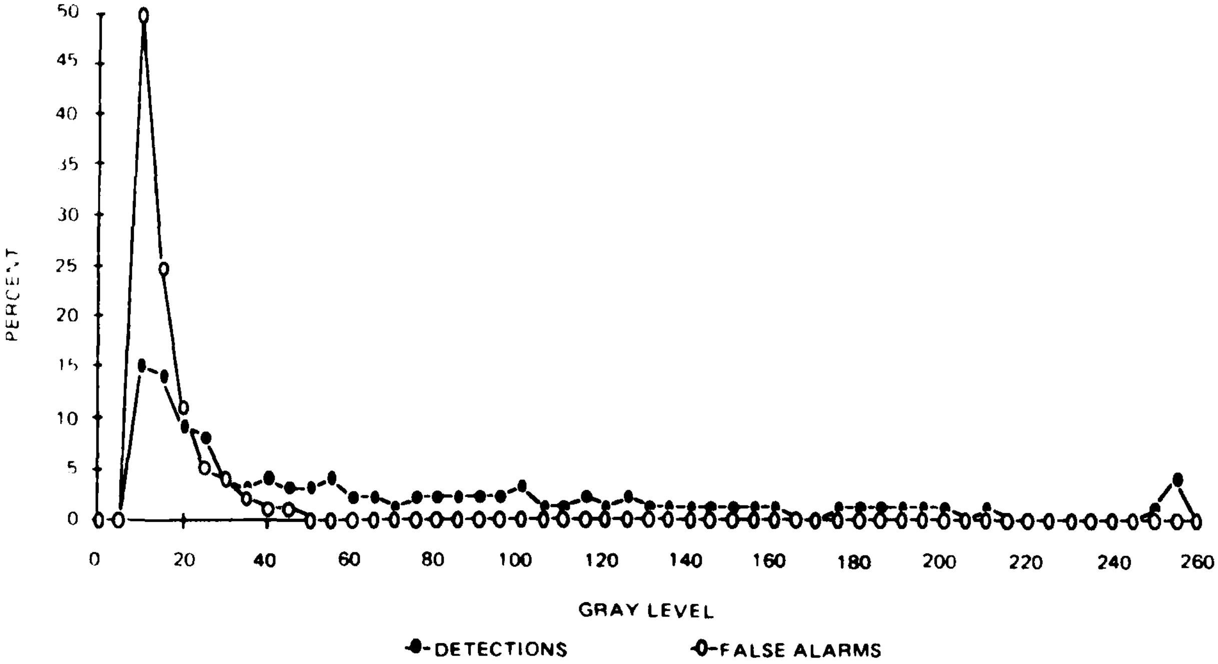

R87-3941-017(T)

Fig. 17 Gray Level Density Function (CP Intensity)

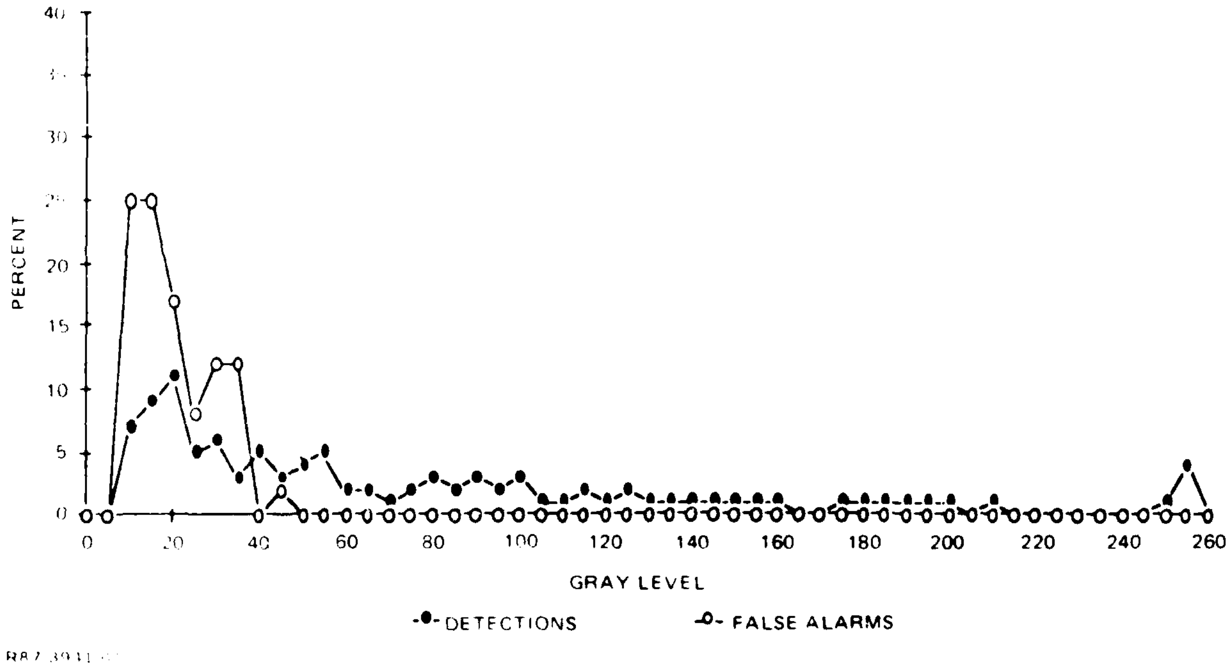

Fig. 18 Gray Level Density Function (CP Shape)

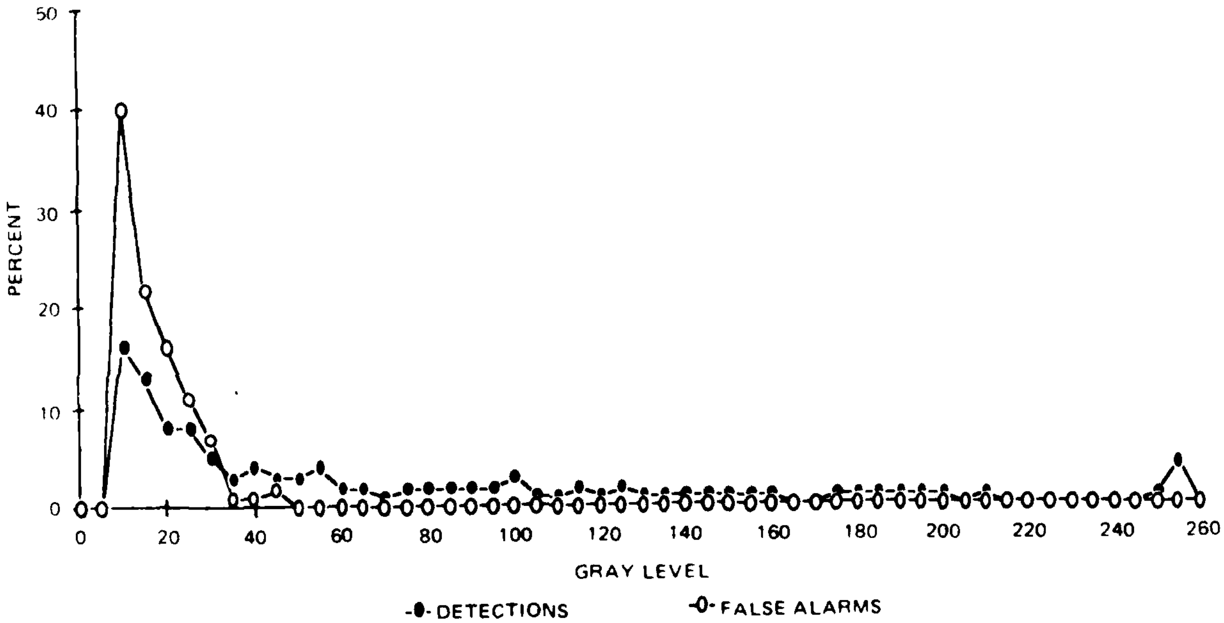

R87-3941.019(T)

Fig. 19 Gray Level Density Function (CP Image Plane)

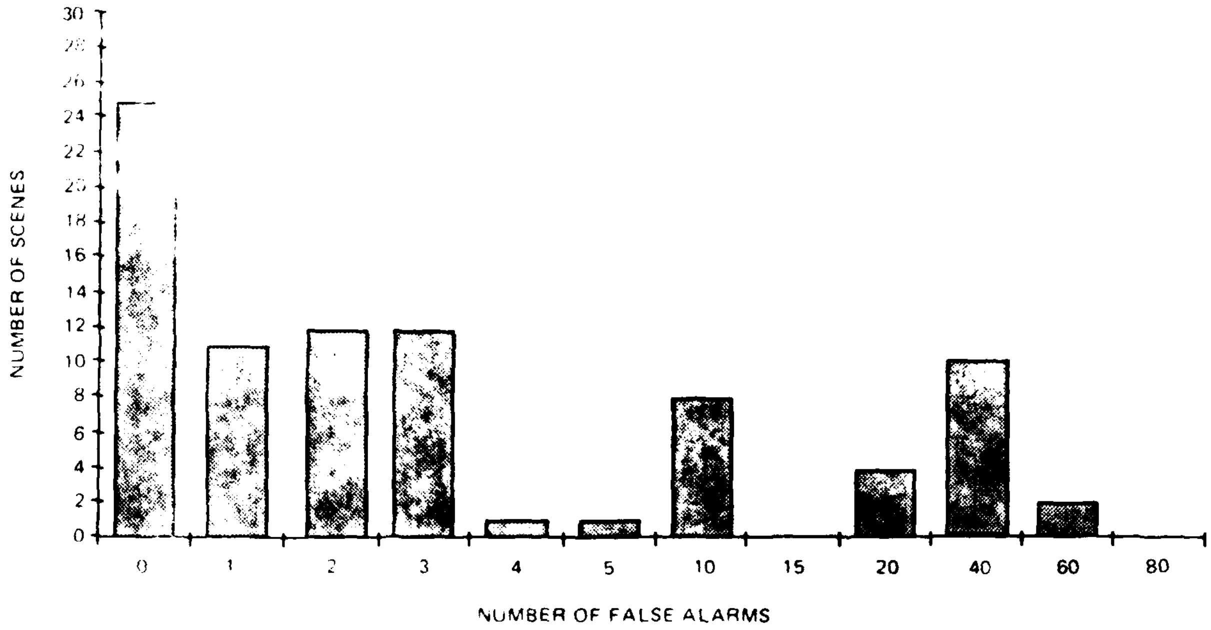

Fig. 20 False Alarm Distribution (CP Intensity)

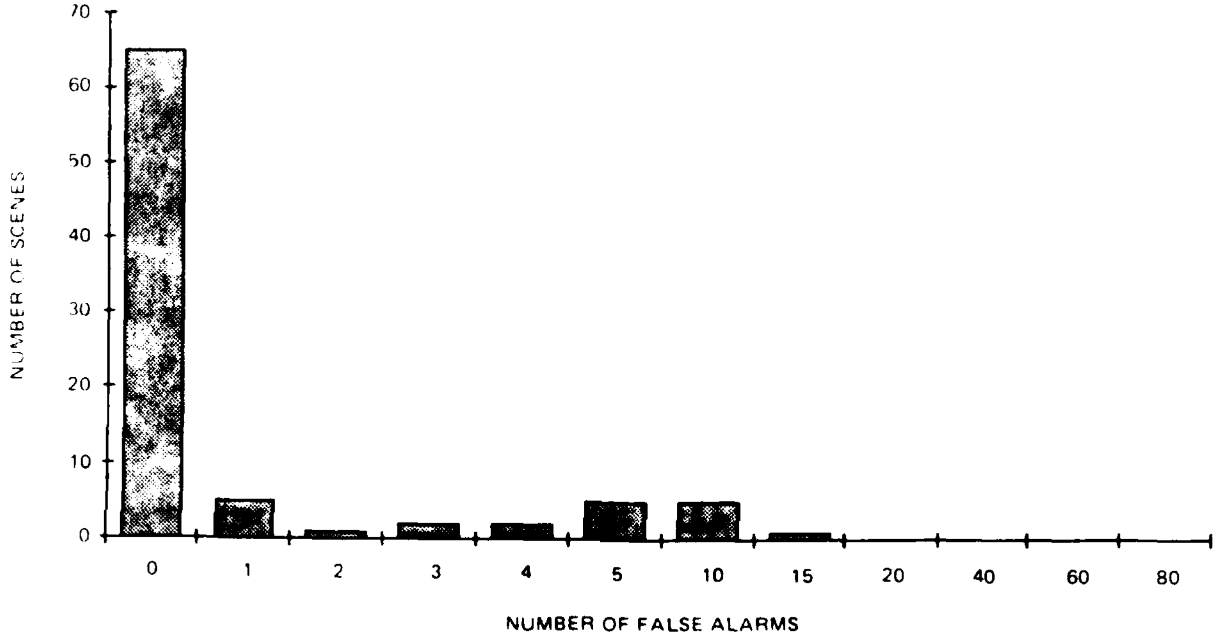

R87-3941-021(T)

Fig. 21 False Alarm Distribution (CP Image Plane)

5.1.8 Target Prioritization on the Basis of Correlation Magnitudes

In Fig. 22-23 we indicate by prioritization of our target choices how strongly ... correlation plane results are related to actual target detect... For example, in Fig. 22, we see that our first choice of target location is correct 80 times in 86 scenes. The break in both Fig. 22 and 23 at choice 3 occurs because there are no scenes in the ERIM development set with more ... three targets. These figures point out that our technique allows ... detections and might be significantly improved by weighting the de... by the strength of the correlation point return.

5.1.9 Signal to Clutter Improvement in the Correlation Plane

In Table 16 we calculated for one example the processing gain that we obtain in the correlation plane in the following way: First, we computed the average gray level in the image and correlation plane. It is clear that this is greatly reduced for the correlation plane. This is partially because the matched filters are made to pass only high-frequency energy. Then we compared the maximum gray level on the target in the scene to the average gray level. We see that this ratio is almost 25 times larger in the correlation plane than in the image plane.

5.1.10 Results for Different Target Types

Table 17 summarizes the results on a target by target basis. We do not know at this point why we do better on some targets than others. It is appropriate to point out at this time that there appears to be little classification possibilities in these results. This is not so surprising in light of the fact that we appear to be detecting high intensity rather than shape. However, if we restrict ourselves to classifying only when the return in the correlation plane is above a given level (i.e., only to the strongest returns) a preliminary survey shows that in these cases we classify correctly at far greater rates than could be attributed to chance. Tables 18-23 show how our detection results are distributed over the various ERIM image metrics. It is intuitively clear from these tables that range, pixels on target, and mean intensity have significant effects on detection and false alarms.

5.1.11 Comparison to Digital Methods

Figures 24 and 25 show the optical results plotted against various

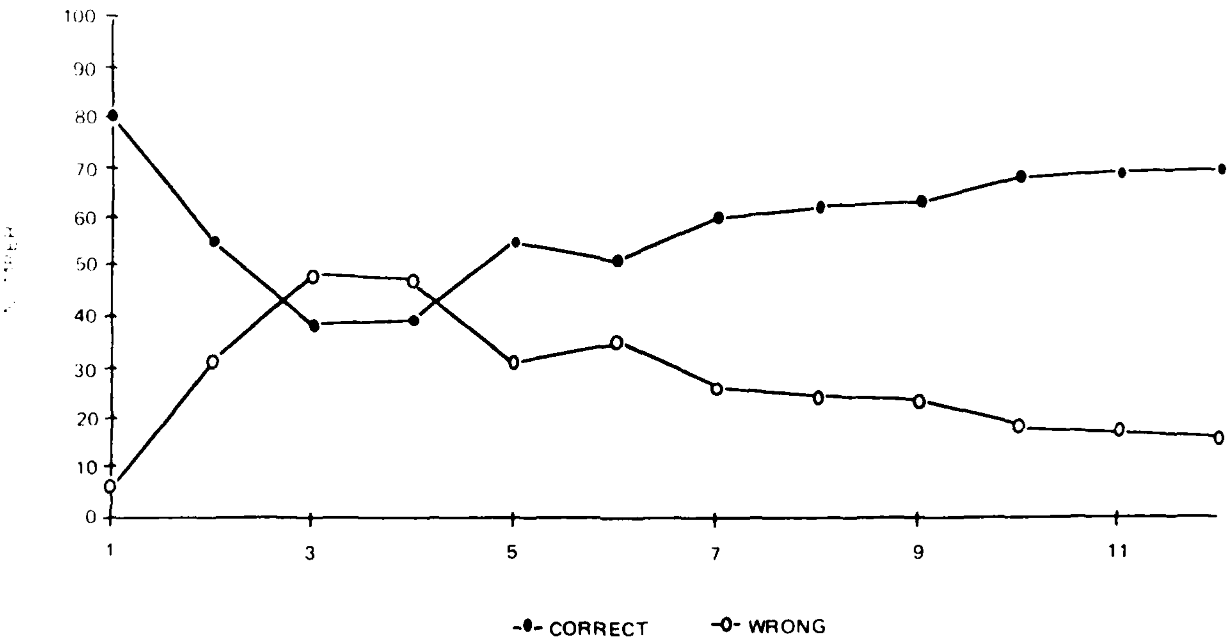

R87-3941-022(T)

Fig. 22 Target Prioritization (CP Intensity)

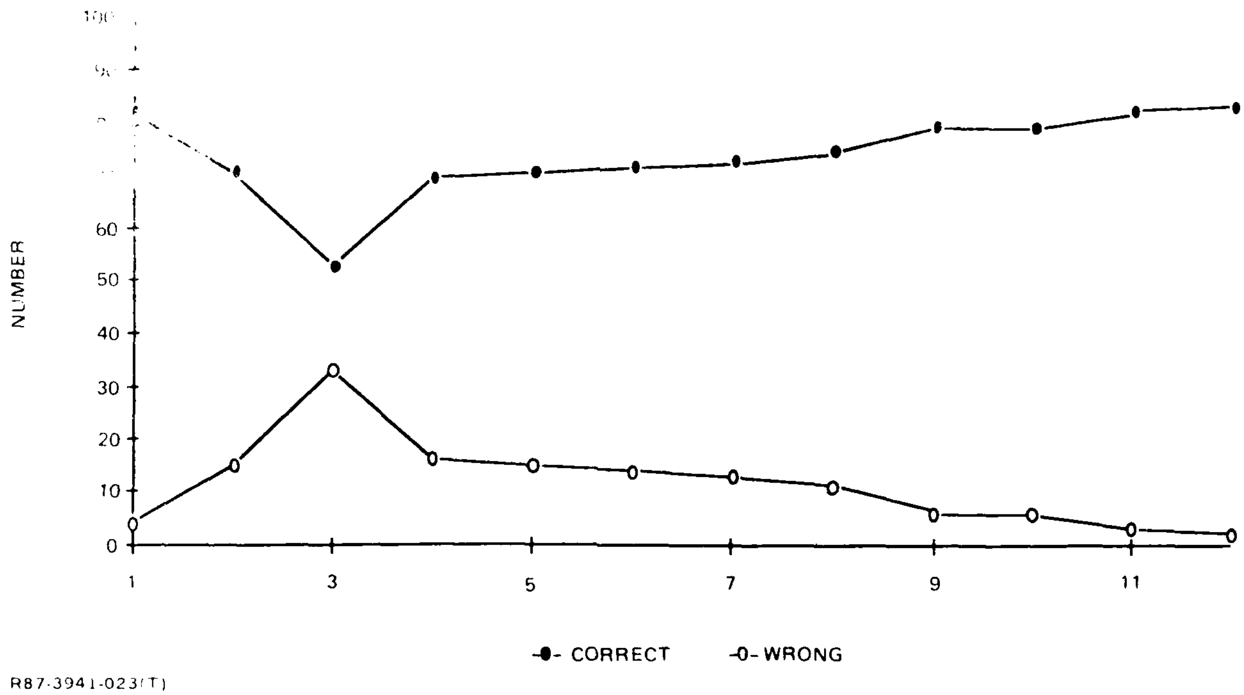

Fig. 23 Target Prioritization (CP Image Plane)

Table 16 Processing Gain Example

SCENE 2.7

	MEAN	SIGMA	MAX	MAX/MEAN
SCENE	56	15	255	4.5
CORRELATION PLANE	1.4	3.1	142	101

R87-3941-045(T)

Table 17 Summary by Targets

TARGET	NUMBER	DETECTED
M-60	9	8
T-55	27	17
M-113	16	14
M-151	17	11
ZIL	18	17
TAB	20	20
APC	25	23
BTR	23	22
UAZ	21	14
M-109	16	13
BRDM	16	12
T-62	6	5
TR2	15	9
BMP	4	3
TOTAL	233	188

R87-3941-046(T)

Table 18 Range Distribution

RANGE	DETECTED	MISSED
0–250	0	0
251–500	0	0
501–750	1	0
751–1000	10	0
1001–1250	25	2
1251–1500	3	1
1501–1750	0	0
1751–2000	21	3
2001–2250	22	4
2251–2500	21	4
2501–2750	31	6
2751–3000	10	1
3001–3250	0	0
3251–3500	22	10
3501–3750	18	7
3751–4750	0	0
4751–5000	2	3
5001–5250	2	4
5251–	0	0

R87-3941-047(T)

Table 19 ESR Distribution

ESR	DETECTED	MISSED
0–25	4	7
26–50	21	7
51–75	139	4
76–500	0	27
501–	24	0
TOTAL	188	45

R87-3941-048(T)

Table 20 TIR2 Distribution

TIR2	DETECTIONS	MISSES
0–5	40	21
6–10	40	14
11–15	29	2
16–20	9	6
21–25	13	2
26–30	11	0
31–35	10	0
36–40	8	0
41–45	5	0
46–50	3	0
51–	20	0
TOTAL	188	45

R87-3941-049(T)

Table 21 TBIR2 Distribution

TBIR2	DETECTIONS	MISSES
0–1	17	7
1–2	35	13
2–3	40	9
3–4	30	5
4–5	21	2
5–6	12	3
6–7	8	2
7–8	7	1
8–9	7	1
9–10	1	2
11–	10	0
TOTAL	188	45

R87-3941-050(T)

Table 22 Pixel Distribution

PIXELS	DETECTIONS	MISSES
0–50	2	5
51–100	13	9
101–150	18	9
151–200	21	3
201–250	19	2
251–300	16	4
301–350	13	2
351–400	10	4
401–450	13	0
451–500	8	0
501–550	4	2
551–600	6	0
601–650	3	0
651–700	1	1
701–750	2	1
751–1000	5	0
1001–	34	2
TOTAL	188	45

R87-3941-051(T)

Table 23 Mean Intensity Distribution

INTENSITY	DETECTIONS	MISSES
0–45	0	0
46–50	0	1
51–55	0	2
56–60	6	7
61–65	21	18
66–70	22	9
71–75	40	2
76–80	36	6
81–85	23	0
86–90	13	0
91–95	9	0
96–100	9	0
101–	9	0
TOTAL	188	45

R87-3941-052(T)

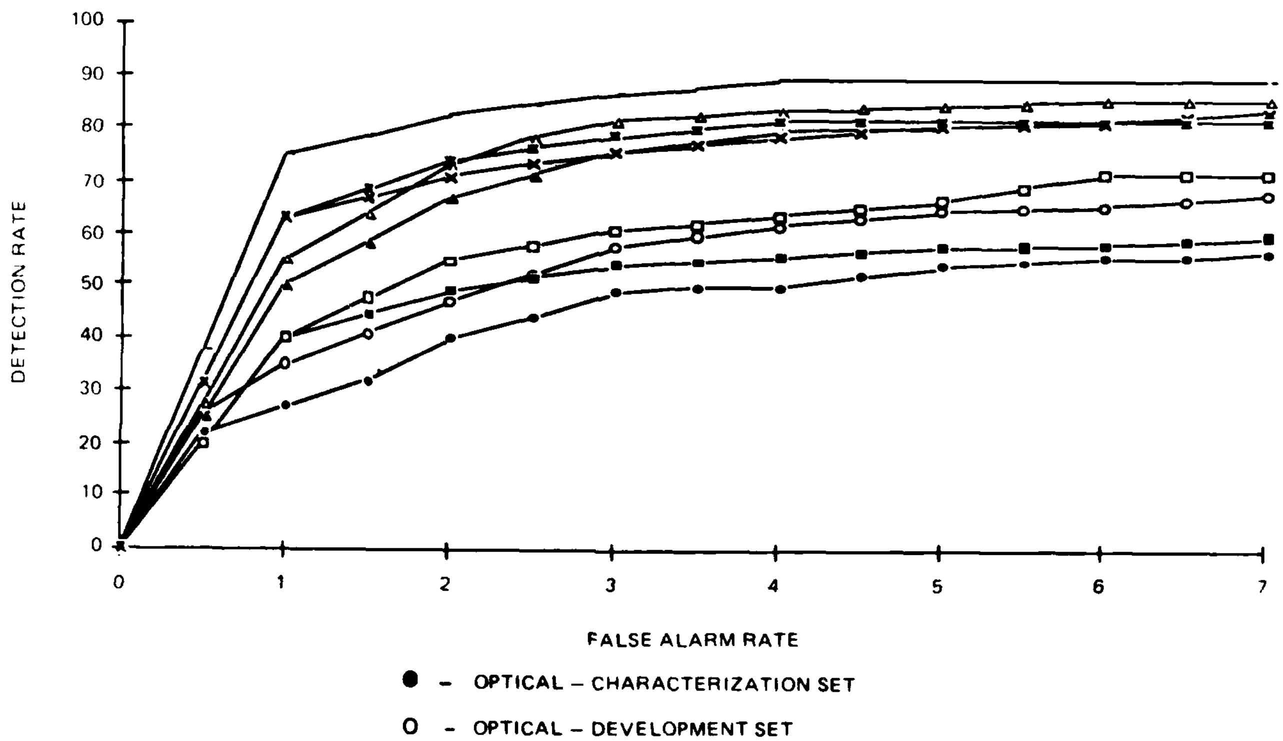

R87-3941-024(T)

Fig. 24 Comparison of Optical vs Various Digital Methods (APG Images – All Ranges)

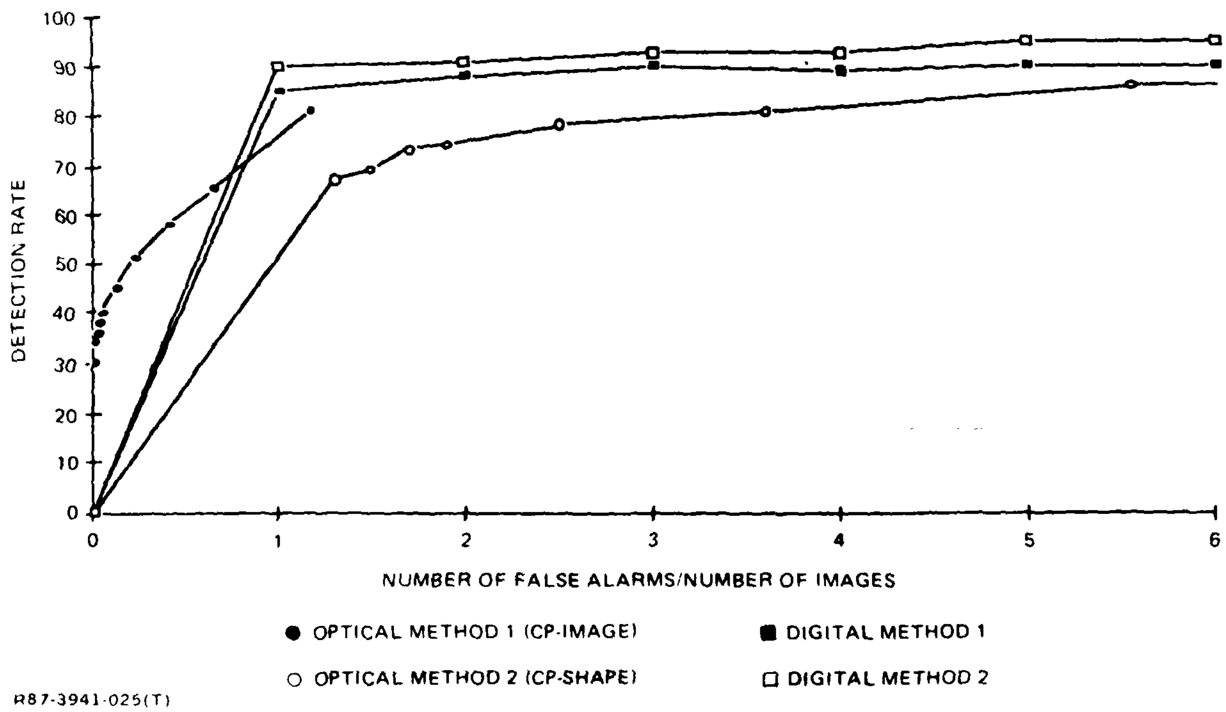

Fig. 25 Comparison of 2 Optical vs 2 Digital Methods; ERIM Development Set (APG Images – Ranges 0-3750)

digital methods. In Fig. 24 ERIM has compiled the characterization set results, over all ranges, for various techniques. We have plotted our development set results with these so that the complete picture could be shown and properly discussed. Figure 24 contains the results of two of the strongest digital methods compared to the optical method for ranges less than 3,750 m. We include this because we feel that the comparison at shorter ranges is more valid; as discussed in this report, the longer range targets give little opportunity for shape matching.

These comparisons must be viewed in the proper context. Several factors should be taken into account. First, we must remember that this is the initial effort for optical correlators against a dataset of IR tactical targets as varied as the ERIM development and characterization sets. Thus, much of the software approach (correlation plane algorithms, matched filter construction methods and synthesis techniques for the multiple matched filters) was under development at the same time the overall technique was being subject to evaluation. After setting certain parametric values, such as the level of low frequency cutoff for the filters and fixing the approaches to the previously mentioned problems, it was not possible to iterate on these values and approaches because of the computational time involved in the simulation. In Section 6 we suggest several important areas where we feel major gains can be made in the improvement of the current results and ways in which the computational problems associated with simulating multiple matched filters can be surmounted. It is also true that the digital algorithms are in a far more mature stage of development than the optical.

Secondly, we were not able to obtain a set of target images which were uniformly distributed in space and from which the matched filter set could be made. Therefore, as we discussed, we used the images which were available, in what we thought would be the best possible way. However, in most cases we were using images to construct matched filters, whose orientations, aspects and ranges were not sufficiently close to those they were tested against in the development and characterization sets. In the context of these handicaps to this specific application, the comparison of the optical methods to the mature digital algorithms is credible and certainly improvable, especially at the shorter ranges.

5.1.12 Comparison of Development and Characterization

In Fig. 24, we notice that the development set results are uniformly better than the results for the characterization set. Since the identical algorithm was applied to both of these datasets some explanation is appropriate. The matched filters that we selected to form the matched filter reference memory were selected from the development set. These same matched filters were then applied to the characterization set. It would appear that if the targets in the two sets were not equivalent from an imaging point of view or statistically equivalent our results would be biased towards the set from which the matched filters were made. Thus, if when evaluating the characterization set we had used matched filters made from that set, we feel, that we would have achieved equivalent results. Finally, if the development set would have contained a sufficient number of target views to realistically cover the views that occurred in the characterization set then we feel that the observed differences would have been minimal.

5.1.13 Conclusions

We believe that these results reinforce the concept that, although infrared targets are for the most part not so well shaped as visual targets and are more dependent on the optical than the geometrical characteristics of the targets, we can use optical techniques to detect and certainly to cue these targets. Further, we conclude, that this set of results has pointed to clear cut areas where improvements could be made in the construction of the individual matched filters and in the composition of the collection so as to raise the detection probability obtained and cut into the false alarm rate.

5.2 WIRE FRAME MATCHED FILTERS

In Fig. 26 the performance of the wire frame matched filters is summarized in operating characteristic curves. In these curves, detection and false alarms levels are related to both range and threshold level. The threshold levels in this portion of the study were based on the standard deviation of the values in the correlation plane. For example, a threshold level of 2, means that all values in the correlation plane above the mean+2 standard deviations were considered as possible target sources. Thus, at high thresholds (4, for example) there were relatively few false alarms, but the detection rate was low. On the opposite end, at a threshold of 1, both

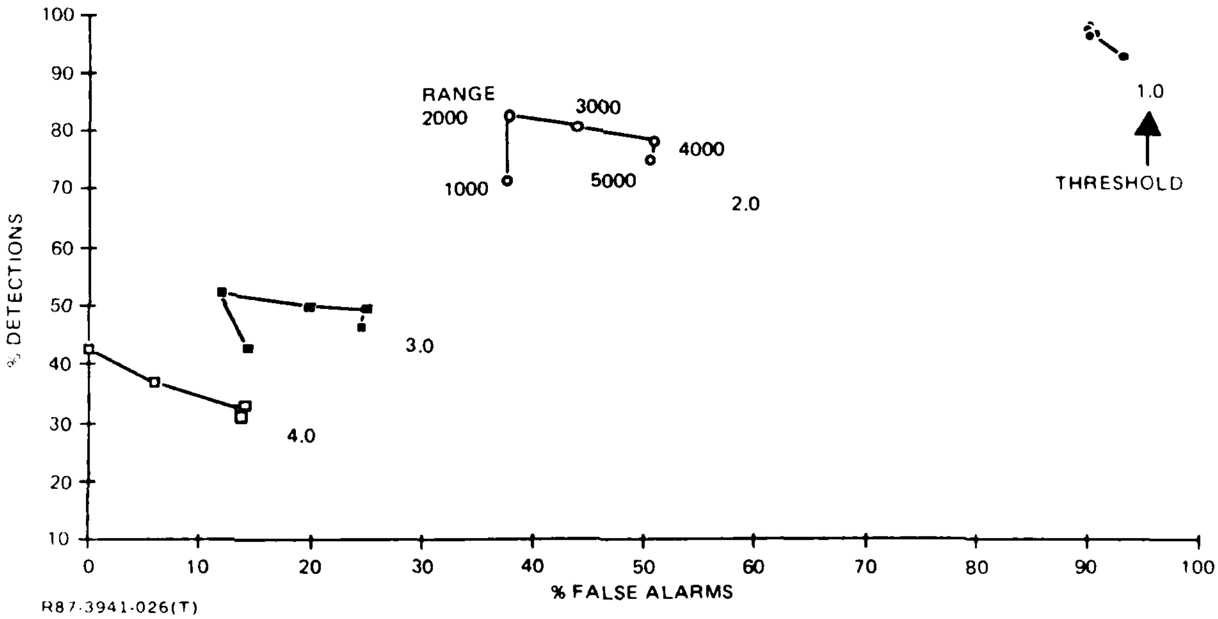

Fig. 26 Operating Characteristics vs Range & Threshold (Wire Frames)

detections and false alarms were very high. At ranges less than 3000 m and for a threshold level of 2, we obtained over 80% detections for a false alarm rate of approximately 42%. The false alarm rate is the ratio of the number of non-targets reported to the total number of targets reported.

5.2.1 Mixing Real and Synthetic Filters

In order to determine if the synthetic filters could be used in conjunction with the real filters (i.e., to fill in the gaps in the real filters) we made a set of matched filters on a 10 deg orientation, 100 m range grid and added in the real filters wherever they existed. We then used this set in exactly the same way as the wire frames only. We found essentially no difference in the results, as can be seen in Fig. 27. The reason for this appears to be that the reals are too sparsely distributed over this grid to have a significant influence on the results. In almost every case, as will be shown, the reals plus wire frames performed identically to the wire frames alone.

5.2.2 Significance of the Density of Synthetic Filters

Finally, to determine the effect of the resolution of the grid on these results we used wire frames at the exact geometries for the targets in the scenes. In this case only one matched filter for each target was correlated against the incoming scene for a total of 14 targets. This set of matched filters represents an upper boundary on what can be accomplished with the wire frames as we have used them. In Fig. 28 the results show only a slight improvement over the wire frame on the 10 deg, 100 m grid. If the wire frames can be improved upon, that improvement probably will come from using better, more thoroughly representative models than those in the ERIM dataset, rather than from more specific knowledge of the scene variables. Further improvements also might be available in better matched filter normalization techniques for size and intensity. There was insufficient time to examine all the possible combinations of filters that might be tried. Some recommendations for studies that might serve to shed some further illumination on the use of synthetic imagery will be suggested in Section 6.

5.2.3 Sensitivity of Wire Frame Results to ERIM Metrics

In Tables 24-30 the results for the wire frames are given in terms of the ERIM metrics. For comparison purposes we also included in these tables the

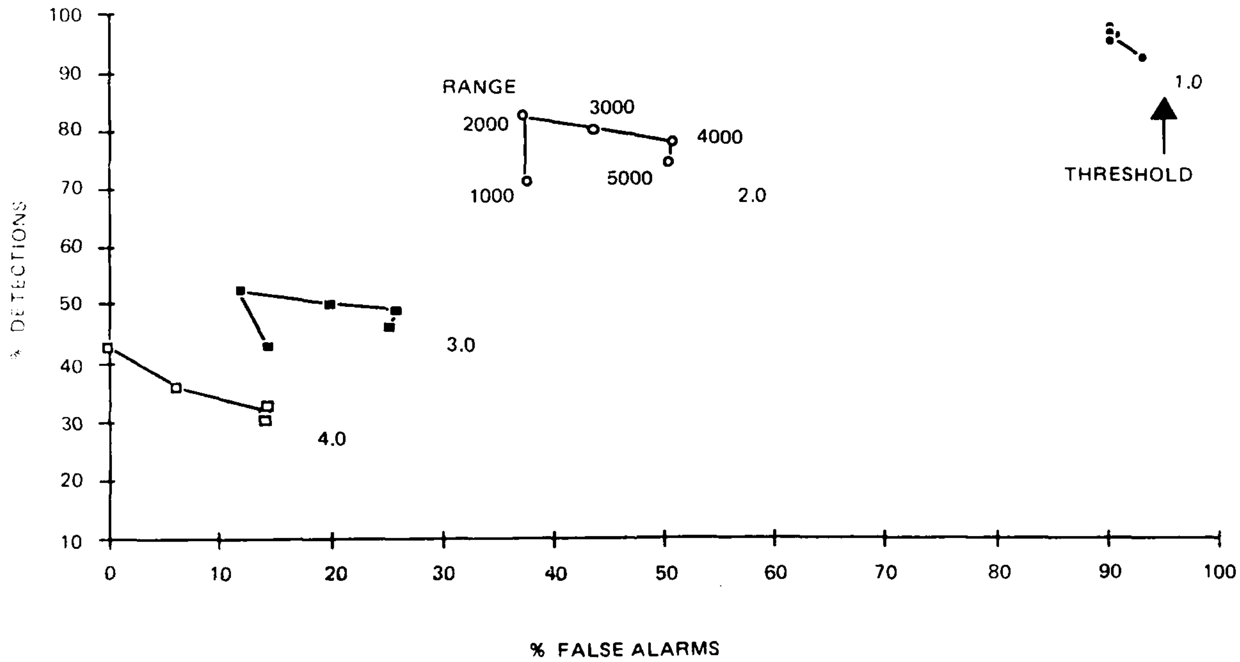

R87-3941-027(T)

Fig. 27 Operating Characteristics vs Range & Threshold (Wire Frames + Reals)

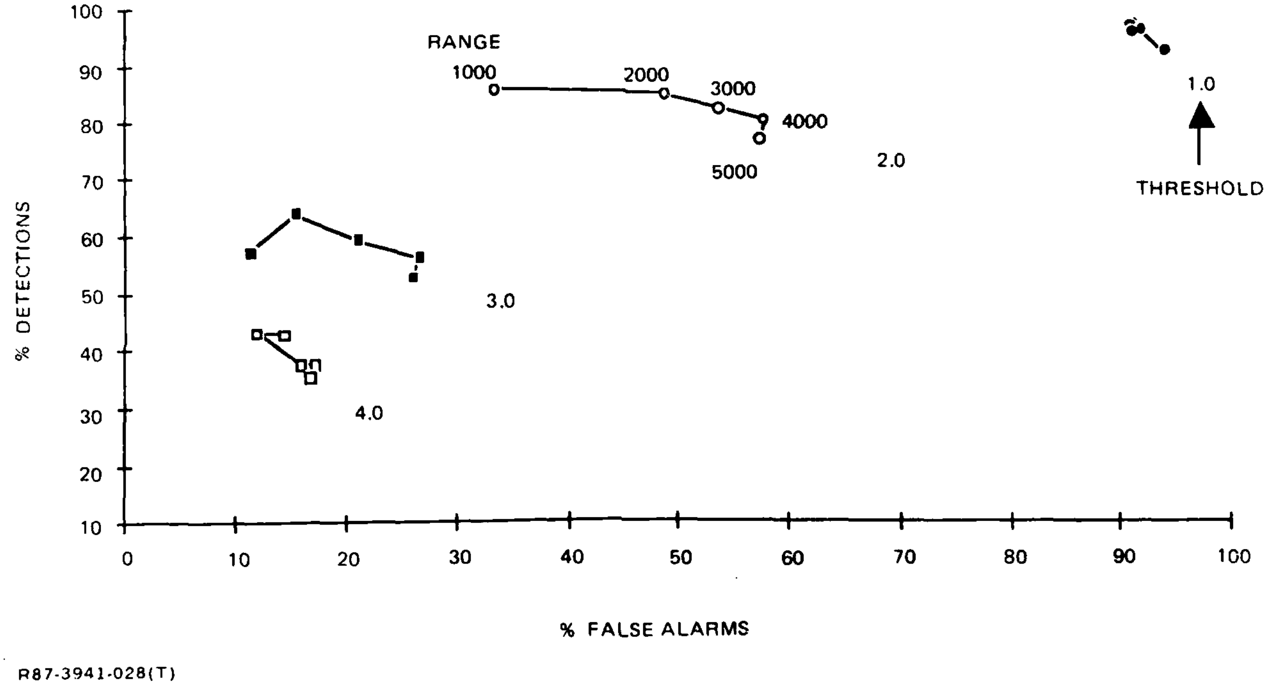

Fig. 28 Operating Characteristics vs Range & Threshold (Exact Wire Frames)

Table 24 Summary by Targets

TARGET	% DETECTED REAL	WIRE FRAME
M-60	89	92
T-55	63	74
M-113	88	71
M-151	65	56
ZIL	94	90
TAB	100	100
APC	92	79
BTR	96	86
UAZ	67	61
M-109	81	83
BRDM	75	50
T-62	83	67
TR2	60	53
BMP	75	50
TOTAL	81	75

R87-3941-053(T)

Table 25 Range Distribution

RANGE	% DETECTED REAL	WIRE FRAME
0–500	0	0
501–1000	85	75
1000–1500	90	90
1501–2000	88	80
2001–2500	84	76
2501–3000	85	82
3001–3500	69	71
3501–4000	72	75
4001–4500	40	–
4501–5001	33	47

R87-3941-054(T)

Table 26 Pixel Distribution

PIXELS	% DETECTED REAL	WIRE FRAME
0–100	52	53
101–200	76	76
201–300	85	81
301–400	79	83
401–500	100	81
501–600	77	89
601–700	80	100
701–1000	88	67
1001–	94	85

R87-3941-055(T)

Table 27 Mean Intensity Distribution

INTENSITY	% DETECTED REAL	WIRE FRAME
0–45	0	0
46–50	0	0
51–55	0	0
56–60	46	38
61–65	54	13
66–70	71	61
71–75	95	77
76–80	86	86
81–85	100	95
86–90	100	100
91–95	100	100
96–100	100	100
101–	100	100

R87-3941-056(T)

Table 28 TBIR2 Distribution

	% DETECTED	
TBIR2	REAL	WIRE FRAME
0–1	71	54
1–2	73	50
2–3	82	68
3–4	86	62
4–5	91	83
5–6	80	95
6–7	80	93
7–8	88	81
8–9	88	91
9–	85	88

R87-3941-057(T)

Table 29 TIR2 Distribution

TIR2	% DETECTED REAL	% DETECTED WIRE FRAME
0–5	66	48
6–10	74	51
11–15	94	80
16–20	60	78
21–25	87	82
26–30	100	91
31–35	100	83
36–40	100	100
41–45	100	100
46–50	100	100
51–	100	95

R87-3941-058(T)

Table 30 ESR Distribution

	% DETECTED	
ESR	REAL	WIRE FRAME
0–25	36	46
26–50	75	48
51–75	97	96
76–500	0	0
501–	100	93

R87-3941-059(T)

results for the real filters over the same range bins (less than 5,200 m).

Table 24 summarizes the overall results and shows that the real filters did slightly better, 81% to 75%, than the wire frames.

In Table 25, the effect of range is clearly shown for both sets of matched filters. There is some indication that the wire frames do slightly better at the longer ranges, but more data would be required before a hard conclusion could be drawn. In Table 26, we see a general tendency to improve as the number of pixels on target increases. This is of course, closely related to range. The strong effect of target intensity is clearly demonstrated in Table 27. For mean intensity on targets greater than 80 the real filters achieve 100% detection and the wire frames do almost as well. Table 29 shows the equivalent fact for TIR2 which is closely related to mean intensity.

5.2.4 <u>Analysis of False Alarms</u>

Rather than report false alarms in terms of the total results, if we examine scene by scene results, then over the 90 scenes including range bins less than 5,000 m, our false alarm rate was approximately 2 false alarms per scene. In Fig. 29 the histograms of false alarms are shown for both real and wire frame filters. The extremely skewed nature of this histogram shows that in the great majority of the scene we reported zero false alarms. The statistics are corrupted by a relatively few scenes in which the filters performed poorly. As mentioned previously, it would seem that a strategy that adapts to the scene contents (i.e., an additional check of some sort when the number of cued targets appears excessive) could reduce the overall false alarm problem.

5.2.5 <u>Conclusions</u>

Great advantages can be had by using synthetic targets for the construction of the reference memory of the optical filter. Certainly, the availability of imagery at any desired range, aspect, or orientation stands out. These results indicate that, if the synthetic imagery can be rendered in greater detail and with further study on how to use these images in the matched filter reference memory, results might be brought to a level that could make synthetic references an interesting possibility.

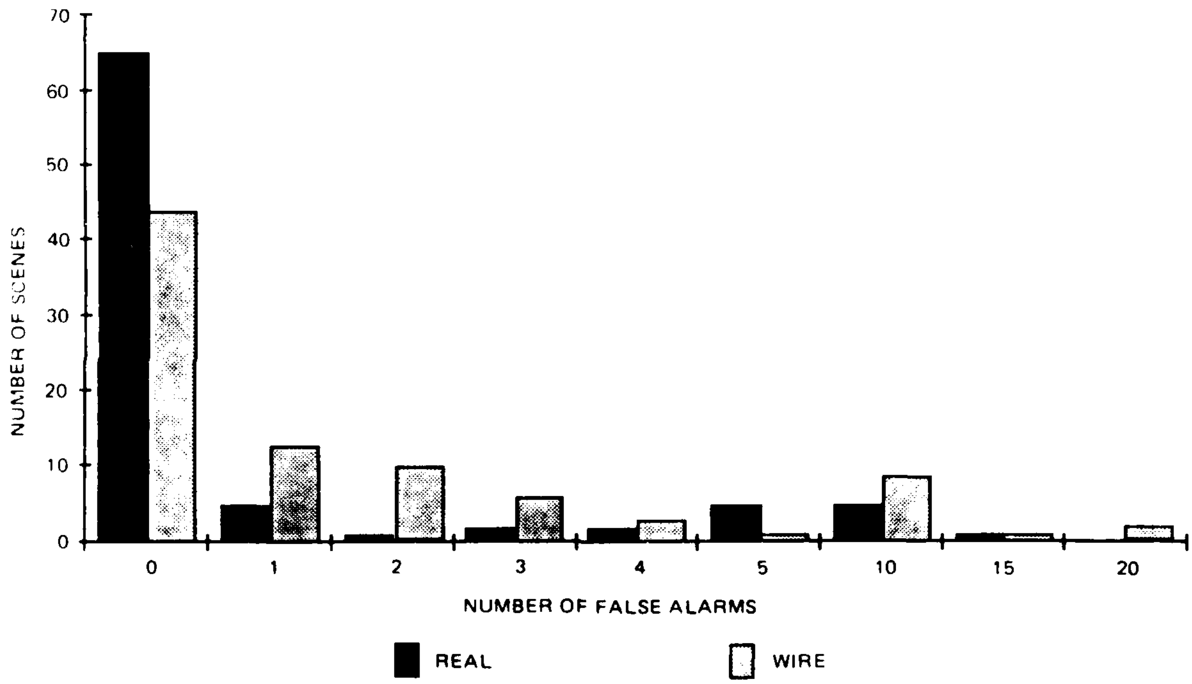

R87-3941-029(T)

Fig. 29 False Alarm Distribution (Real vs Wire)

6.0 CONCLUSIONS & RECOMMENDATIONS

The principal conclusions of this study are the following:

1. Using images from the actual dataset to construct the matched filter reference memory, we were able to obtain approximately 82% detections with a false alarm rate of 1.1/scene, for targets at less than 5,200 m. These results were obtained despite the fact that a systematically placed sample of real images was not available during the course of the study.

2. Using synthetically generated reference images, we were able to obtain results of 78% detections with a false alarm rate of 2/scene. These results could be greatly improved using more detailed synthetic imagery.

3. In the IR imagery we studied, the matched filters were much more sensitive to intensity than they were to shape which explains the high false alarm rates that existed in some individual scenes.

4. Lack of time prevented in-depth investigation and analysis of classification capabilities but initial results are poor. This is consistent with the observation that the filters are more sensitive to intensity than to shape. Significant improvement could be possible if better filter normalization techniques are developed.

5. The principal areas that should be investigated for improving the results above are correlation plane analysis and matched filter normalizations. New techniques for carrying these functions could result in significant improvement. Detailed suggestions for future efforts which would be aimed at accomplishing improved results are given in the following section.

As a result of the CAPIR study, we now have our first indications of how Optical Matched Filtering performs on a real world infrared image set and how it compares to some of the digital techniques. These comparisons should be tempered by the fact that this is the first time this approach has been attempted on as diverse and shapeless imagery as was available in the ERIM development and characterization sets. We believe strongly that these efforts

represent a foundation from which significant improvements can be made. During the past year, we also developed a great deal of the simulation software as it applies to the CRAY 1M, and the future program will benefit from having all of that essential work completed. Finally, in future studies we will have the availability of a newer, faster, and bigger (from a memory point of view) CRAY, the CRAY XMP. This will also benefit future programs which, like CAPIR, have extraordinary computational demands.

The purpose of the following is to suggest, on the basis of the CAPIR experiences, some future directions for the research, which we think will be most beneficial to the ultimate development of an Optical Matched Filter for target detection.

6.0.1 Improvement of Current Results

We believe that we can improve our current results with respect to detection, false alarms and classification. Several approaches to achieving this are possible:

6.0.2 Matched Filter Construction Methods

An investigation of matched filter construction techniques including the design of improved high pass filters and methods for energy normalization could prove fruitful. One approach to this could be to use optical versions of standard digital edge enhancers including some of those reported on by other CAPIR contractors

6.0.3 Energy Normalization Methods

Energy normalization which we began to use at the end of the contract in our correlation plane analysis could help improve classification results. These should include ways of accounting for the differences in total energy in the filters due to size differences, and for matched filters which are constructed from real images due to contrast differences. Also a method to account for local contrast and intensity differences in the image plane needs to be investigated.

6.0.4 Correlation Plane Analysis Methods

We should try new methods of correlation plane (CP) analysis which take into account more than just CP intensity. Other variables such as size and shape should be taken into account. Another thing that should be tried is to apply some of the results of the other CAPIR studies on our CP. The improved

signal-to-clutter ratio that we gain in the CP combined with the techniques that have been developed by other CAPIR contractors for image plane target detection should allow us to achieve an improvement especially with respect to false alarm rates. Of special interest is the method presented in Ref 2 in which a preprocessor extracts the statistics of contrast, border strength and hot spots. This approach could be used in its entirety to do the correlation plane analysis or could be used as a preprocessor to determine if the hot spot is dominant, in which case, correlation may not be appropriate and a different technique could be used.

6.0.5 Further Analysis of Autocorrelations in the Current Sample

It would also be beneficial to spend some time trying to get a better understanding of why some of the autocorrelations are not producing peaks of sufficient intensity to either detect or classify their target and how this could be corrected. In focusing on autocorrelations we can eliminate many variables which complicate the understanding of the results and perhaps see more clearly which of the above techniques can yield the most significant gains.

6.0.6 Comparisons of Multiple Matched Filter to Other Optical Algorithms

We should compare our approach to other typical optical approaches using the ERIM datasets as a basis for comparison. Among the techniques that should be investigated are synthetic discriminant functions, circular harmonics (size and rotation invariant), and binary phase only. We believe that having a common comparison of all of these techniques would be a great benefit to the Optical Correlator community.

6.0.7 Classification

One of the advantages of the optical approach is that it should combine detection and classification in one process. However, we were not seeing much classification capability in our early results. This could be due to the lack of shape in the targets or it could be because of the fact that the targets are of such different sizes that the larger ones tend to collect more energy than the smaller ones. If the latter is true, then correlation plane energy normalization could be effective in improving the results.

In our future efforts we plan to extract a small subset of the ERIM development set, perhaps 20-25 images, on which we will try out and test ideas

AD-A193 059 OPTICAL MATCHED FILTERS FOR AUTONOMOUS INFRARED SEEKERS 2/2
(U) GRUMMAN CORP BETHPAGE NY SYSTEMS SCIENCES
DIRECTORATE J MENDELSOHN ET AL. 14 MAY 87 RE-735

UNCLASSIFIED NV-1-13 DAAL02-85-C-0144 F/G 17/5.1 NL

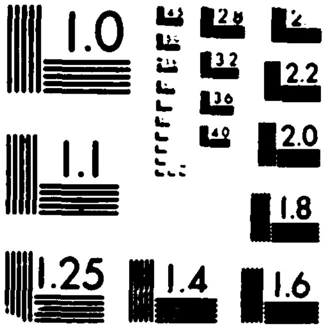

MICROCOPY RESOLUTION TEST CHART

before applying them to the complete development or characterization sets. Since we are now well aware of which images are easy and which cause problems, we can be more efficient in our approach. Thus, by using the experience gained in this effort we will be able to study methods for improving the results in a more efficient and computationally less intensive environment.

7.0 ACKNOWLEDGMENT

The authors would like to acknowledge the support and encouragement of Byong Ahn of the Center for Night Vision and Electrooptics. In the DOD community, he has been, and continues to be, one of the main supporters of optical techniques in target recognition.

8.0 REFERENCES

1. Mendelsohn, J., "Analysis of IR Imagery by Optical Correlation Techniques," Grumman CRC Report RE-687, Final Report on Contract DAAK20-83-C-0162, Aug 1984.

2. Schachter, B., et al., "Analysis and Evaluation of IR Target Detection, Segmentation and Classification Algorithms," Westinghouse Electric Corp., CAPIR Review, Dec 9-10, 1986.

END
DATE
FILMED
8-88
DTIC